



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 528 455 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
04.05.2005 Bulletin 2005/18

(51) Int Cl.7: **G06F 1/00**

(21) Application number: **04025826.1**

(22) Date of filing: **29.10.2004**

(84) Designated Contracting States:
**AT BE BG CH CY CZ DE DK EE ES FI FR GB GR
HU IE IT LI LU MC NL PL PT RO SE SI SK TR**
Designated Extension States:
AL HR LT LV MK

(72) Inventors:
• **Shapiro, William M.**
Palo Alto CA 94301 (US)
• **Herbach, Jonathan D.**
Mountain View CA 94043-4923 (US)
• **Donahue, James**
Sunnyvale CA 94087 (US)

(30) Priority: **31.10.2003 US 699124**

(71) Applicant: **Adobe Systems Incorporated**
San Jose, California 95110 (US)

(74) Representative: **Müller-Boré & Partner**
Patentanwälte
Grafinger Strasse 2
81671 München (DE)

(54) **Offline access in a document control system**

(57) Systems and techniques to provide offline access in a document control system. In general, in one implementation, the technique includes: receiving a request from a client, and pre-authorizing the client, in response to the request, to allow actions by a user as a member of a group of users by sending to the client offline access information including a first key associated with the group, the first key being useable at the client

to access an electronic document by decrypting a second key in the electronic document. Receiving a request can involve receiving a request from the client to take an action with respect to a second document. The technique can also include verifying the user at the client as an authenticated user, and the offline access information can include user-specific keys, group-specific keys, a policy, and a document revocation list.




FIG. 1

EP 1 528 455 A1

Description

BACKGROUND OF THE INVENTION

[0001] The present application describes systems and techniques relating to document control, for example, offline access control in a document control system.

[0002] Traditional document control systems have included servers that store and manage encryption keys for documents secured by the system, providing persistent protection for documents by requiring the server to be contacted before a secured document can be opened. Such systems have also provided offline capabilities by caching a cryptographic document key on a client to allow the client to open a document for a limited time when the user is offline, provided the document is first opened while online. Such systems have also been able to log document access information, including caching of log information while offline, for use in auditing document access.

[0003] Conventional document management systems have included document permissions information associated with documents that allow different groups of individuals to have different permissions, and conventional document viewing software applications have also included software plug-ins designed to translate document permissions information from a document management system format to a format used by the software application, i.e., a separate software plug-in required for each integration with a document management system. Moreover, the eXtensible Rights Markup Language (XrML™) is being defined to theoretically allow a document viewing application to understand resources and permissions from any system that complies with the XrML™ rules.

[0004] Many different encryption schemes have been used to secure documents. These have included symmetric encryption on a per-document basis, requiring individuals to remember passwords for individual documents, and combined asymmetric-symmetric encryption schemes (e.g., Pretty Good Privacy (PGP™) encryption) that provide the ability to decrypt multiple documents based on the user's single password. In the network multicast/broadcast context, various encryption protocols have also been used that cache encryption keys on clients. Many software products directly integrate with existing enterprise authentication systems (e.g., Lightweight Directory Access Protocol). Moreover, various systems have also provided functionality to allow users to find the most recent version of a distributed document, such as the Tumbleweed Messaging Management System™, which secures e-mail systems and can send a recipient of an email with an attached document an email notification when the original version of the attached document is updated, where the email notification has a URL (Universal Resource Locator) link back to the current document.

SUMMARY OF THE INVENTION

[0005] In general, in one aspect, the invention features operations including receiving a request from a client, and pre-authorizing the client, in response to the request, to allow actions by a user as a member of a group of users by sending to the client offline access information including a first key associated with the group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document. Pre-authorizing the client can involve comparing a time of last recorded client-synchronization with a time of last change in user-group information for the user. Pre-authorizing the client can involve comparing current user-group information for the user with received user-group information for the user from the client.

[0006] The client can allow actions with respect to the electronic document based on document-permissions information residing in the electronic document. The offline access information can include document-permissions information associated with multiple documents, including the electronic document, and the client can allow actions with respect to the electronic document based on these document-permissions information.

[0007] Receiving a request can involve receiving a request from the client to take an action with respect to a second document. The operations can also include verifying the user at the client as an authenticated user, and the offline access information can include at least one user-specific key, at least one group-specific key, including the first key, at least one set of document-permissions information associated with multiple documents, and a document revocation list. The at least one set of document-permissions information can be one or more policies associated with the second document. The offline access information can include at least one set of document-permissions information associated with a specific document selected based on synchronization prioritization information. The operations can also include receiving an offline audit log from the client.

[0008] According to another aspect, the invention features operations including receiving from a document control server, when online, offline access information including a first key associated with a group of users of the document control server, and allowing access to an electronic document, when offline, by performing operations including using the first key to decrypt a second key in the electronic document and governing actions with respect to the electronic document based on document-permissions information associated with the electronic document. Governing actions with respect to the electronic document can involve obtaining the document-permissions information from the electronic document. Governing actions with respect to the electronic document can involve identifying a document policy reference in the electronic document, and obtaining locally retained document-permissions information based on

the document policy reference.

[0009] The operations can also include preventing access to the document, when offline, if a difference between a current time and a receipt time of the offline access information exceeds a server-synchronization-frequency parameter. The server-synchronization-frequency parameter can be specific to the document. The operations can also include maintaining an offline audit log, and uploading the offline audit log when online.

[0010] According to another aspect, the invention features operations including encrypting an electronic document, and incorporating into the encrypted electronic document an address of a document control server, document-permissions information, and an encryption key useable in decrypting the encrypted electronic document, the encryption key being encrypted with a key generated by, and associated with a group of users of, the document control server. The encryption key can be a session key generated by the document control server, and encrypting the electronic document can involve encrypting the electronic document using a document key.

[0011] Incorporating can involve incorporating into the encrypted electronic document a document security payload including the document key and the document-permissions information, the document security payload being encrypted using the session key. The document security payload can also include a document identifier assigned by the document control server, and incorporating can also involve incorporating into the encrypted electronic document a copy of the session key encrypted using a public key associated with the document control server. Moreover, the document-permissions information can specify access permissions at a level of granularity smaller than the electronic document itself.

[0012] According to another aspect, a system includes a document control server that synchronizes offline access information with a client in response to a client request, the offline access information including a first key associated with a group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document. The system includes a client that allows access to the electronic document, when offline, by a user as a member of the group, using the first key to decrypt the second key in the electronic document and governing actions with respect to the electronic document based on document-permissions information associated with the electronic document. The client can include an agent that periodically contacts the document control server to synchronize the offline access information.

[0013] The server can include a server core with configuration and logging components, an internal services component that provides functionality across dynamically loaded methods, and dynamically loaded external service providers, including one or more access control service providers. The system can also include a business logic tier having a cluster of document control serv-

ers, including the document control server, an application tier having the client including a viewer client, a securing client, and an administration client, and a load balancer that routes client requests to the document control servers.

[0014] The client request can be a request from the client to take an action with respect to a second document. The document control server can be a permissions-broker server including a translation component. The second document can be a document secured previously by the permissions-broker server, and the translation component can be operable to translate first document-permissions information in a first permissions-definition format into second document-permissions information in a second permissions-definition format in response to the request being received from the client.

[0015] The server can be a permissions-broker server operable to identify information associated with the second document in response to the request. The associated information can be retained at the server and indicate a third electronic document different from and associated with the second document. The server can be operable to relate information concerning the third electronic document to the client to facilitate the action to be taken.

[0016] The server can be a permissions-broker server operable to obtain and send, in response to the request, a software program having instructions operable to cause one or more data processing apparatus to perform operations effecting an authentication procedure. The client can use the authentication program to identify a current user and control the action with respect to the second document based on the current user and document-permissions information associated with the second document.

[0017] The invention can be implemented to realize one or more of the following advantages. A document control system can be easily and tightly integrated with existing enterprise infrastructure, such as document management systems, storage systems, and authentication and access control mechanisms. Users of the document control system can be enabled to perform authorized actions with minimal annoyance. Client installations can be minimized, and server-initiated, transparent client customization can be performed, thereby easing the process of deploying an enterprise solution. Moreover, the document control system can be deployed on multiple platforms and not be intimately tied to a particular platform.

[0018] Functionality of a client can be pushed onto a server, enabling simplified management and deployment by minimizing the size and complexity of the client application. New functionality affecting the client's operations can be implemented at the server without requiring complicated client updates. An authentication system can allow authentication processes to be plugged into a client as needed. The authentication system can be integrated with multiple different authentication

mechanisms, including later developed authentication mechanisms. The authentication system can support transparent authentication, such as by caching a logon ticket for a period of time or by re-authenticating the user transparently. The authentication system can ease deployment. A server administrator can configure user authentication with a newly developed authentication process, and the new authentication process can be automatically plugged into the client when authentication is to occur. The authentication process can be independent of document permissions and actions, and thus authentication can occur in between actions without needing to take the nature of the distribution of existing documents into consideration.

[0019] A document control system can use an existing client application with its own representation of document-permissions information, which may not be able to represent all the document-permissions information used for a document to be controlled. Server plug-ins can translate between document-permissions information types, allowing additional document protection concepts, such as may be present in an enterprise system, to be used with the existing client application. Dynamic translation and specification of document-permissions information at a permissions broker server can create a highly versatile and readily upgradeable document control system. The system can be readily made backwards compatible because the system can translate to older formats of permissions as well as provide a different document altogether if needed.

[0020] An offline access model can be provided in which a user can be offline the first time they access a document. The user need not be prevented from accessing a document that they have permission to access simply because they are offline, while at the same time the user can be prevented from accessing a document offline that they have been denied access to while online. A bounded time can be provided between when a revocation is issued and when it takes effect on all clients in the system. Moreover, a bounded time can be provided between when a policy is modified and when all clients use the modified policy.

[0021] A document information delivery technique can be provided that automatically sends information concerning different versions of a document to be accessed. A document can be tethered to a document control system as described, and when the document is opened, the system can relate information concerning a different document that should be accessed instead. Various workflows can be defined. Such workflows can ensure that users always have the latest version of a document or provide customized user-dependent document delivery. Viewing of a different document can be suggested to or forced on the user, or both can be possible dependent upon the document.

[0022] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features and advan-

tages of the invention will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] FIG. 1 is a block diagram illustrating an operational environment for a document control system.

[0024] FIG. 2 is a block diagram illustrating an example document control server.

[0025] FIG. 3 is a block diagram illustrating workflow in an authentication system.

[0026] FIG. 4 is a flow chart illustrating an authentication technique employed by a server.

[0027] FIG. 5 is a block diagram illustrating workflow in a document control system.

[0028] FIG. 6 is a flow chart illustrating a document control technique employed by a permissions-broker server.

[0029] FIG. 7 is a block diagram illustrating workflow in a document control system integrated with a document repository.

[0030] FIG. 8 is a block diagram illustrating workflow in a document control system integrated with an email client.

[0031] FIG. 9 is a block diagram illustrating a document control server corresponding to the example of FIG. 2.

[0032] FIG. 10 is a block diagram illustrating example details of the server from FIG. 9.

[0033] FIG. 11 is a block diagram illustrating an offline document access model as can be used in a document control system.

[0034] FIG. 12 is a flow chart illustrating a synchronization operation as performed by a server.

[0035] FIG. 13 is a flow chart illustrating a synchronization operation as performed by a client.

[0036] FIG. 14 is a block diagram illustrating components of a secured document.

[0037] FIG. 15 is a flow chart illustrating a document information delivery technique employed by a server.

[0038] FIG. 16 is a block diagram illustrating workflow in a document control system.

[0039] FIG. 17 is a flow chart illustrating a document information receiving technique employed by a client.

[0040] FIG. 18 is a block diagram illustrating document securing workflow in the document control server of FIG. 9.

[0041] FIG. 19 is a block diagram illustrating server-side access control list evaluation workflow in the document control server of FIG. 9.

[0042] FIG. 20 is a block diagram illustrating online document viewing workflow in the document control server of FIG. 9.

[0043] FIG. 21 is a block diagram illustrating revocation workflow in the document control server of FIG. 9.

[0044] FIG. 22 is a block diagram illustrating audit events retrieval workflow in the document control server of FIG. 9.

[0045] FIG. 23 is a block diagram illustrating a document control system with multiple document control servers.

[0046] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0047] The systems and techniques described can be used to realize a document control system, such as may be used by an enterprise in connection with document management. The document control system can operate as a stand-alone system or as a component of another system. The document control system can provide persistent document security by controlling who can view documents and what can be done with them, regardless of where the document resides. As used herein, the terms "document" and "electronic document" mean a set of electronic data, including both electronic data stored in a file and electronic data received over a network, which can be represented as a single document icon in a graphical user interface of an operating system (OS) or software application. An electronic document does not necessarily correspond to a file. A document may be stored in a portion of a file that holds other documents, in a single file dedicated to the document in question, or in a set of coordinated files. Additionally, as used herein, the term "periodically" means recurring from time to time, and does not require regular intervals.

[0048] The systems and techniques described can be used with many different types of documents, including, for example, PORTABLE DOCUMENT FORMAT™ (PDF™) documents. PDF™ documents are in a format originated by Adobe Systems Incorporated of San Jose, California. A PDF™ document is an example of an electronic document in a platform-independent document format that can define an appearance of the electronic document. This document format can be a platform independent storage format capable of storing many different types of data, including graphics, animation and sound, and the defined appearance can be defined for multiple types of display devices, providing a document originator with control over the look and feel of the document regardless of the final destination device. Using documents in this type of format with the techniques described can result in additional advantages for the resulting systems. For example, the document control system can have an architecture that is not tied to a particular software development platform (e.g., the system can be designed to run on both Java and .NET), and can use platform-independent documents, such as PDF™ documents. Thus, the document control system can readily function across several platforms.

[0049] FIG. 1 is a block diagram illustrating an operational environment for a document control system. A network 100 provides communication links between one or more clients 110, one or more servers 120, and one or more enterprise systems 130. The network 100 may

be any communication network linking machines capable of communicating using one or more networking protocols, including a local area network (LAN), metropolitan area network (MAN), wide area network (WAN), enterprise network, virtual private network (VPN), and/or the Internet. A client 110 can be any machine(s) or process(es) capable of communicating over the network 100 with a server 120, and the server 120 can be any machine(s) or process(es) capable of communicating over the network 100 with an enterprise system 130. Moreover, the client(s) 110 can also communicate with the enterprise system(s) 130.

[0050] The enterprise system(s) 130 can be a storage system, an authentication system, a communication system, and/or a document management system. The server(s) 120 can be designed to tightly integrate with existing enterprise system(s) 130 and leverage existing enterprise infrastructure. For example, the server(s) 120 can provide rich support for user and group information in enterprises, where such information may come from multiple sources, as is common in large companies that have been involved in recent mergers. The server(s) 120 can provide document security while being minimally obtrusive, making the system easier to use and thus easier to deploy effectively. For example, the server(s) 120 can implement a document control system that provides a sophisticated offline-access mechanism, as described further below, that allows users to view documents while offline, even if they have not previously viewed the document while online. Thus, the document control system can maintain a low-profile during normal operation, making the presence of document security less visible, and thus more usable.

[0051] FIG. 2 is a block diagram illustrating an example document control server 200. The document control server 200 can include a server core 210 with configuration and logging components 220, 230. The server core 210 can provide a remote procedure call (RPC) interface to the clients that contact the server 200. An internal services component 240 can provide functionality across methods 250. Other components of the server 200, including the methods 250 and external service providers 260, can be dynamically loaded based on information provided by the configuration component 220. The methods 250 can specify the functionality that the server 200 exports to the clients (e.g., secure a document, execute an audit query, etc). The external service providers 260 can specify external facilities that are available to the methods 250 (e.g., storing data, authenticating users, etc).

[0052] The configuration component 220 can define an interface to a configuration object, and the logging component 230 can define an interface to a logging object used by the server 200 to log a wide variety of information. The configuration object can be a server configuration file (e.g., a ".ini" file read by the server 200), and the logging object can be a log file (e.g., a text file). Alternatively, the configuration object and the logging

object can be local or remote objects defined using a standardized interface (e.g., the java standards JMX (java management extension) and log4j, respectively).

[0053] The RPC interface provided by the server core 210 can be used to present a method interface to the clients: a client can RPC each named method and provide an appropriate set of arguments. The server 200 can initialize itself by reading a set of method classes that export the server method interface and define the methods that the server 200 will make available to clients. The internal services 240 can be internal components of the server that are used across all of the methods 250. These can be statically defined and/or dynamically loaded as dependencies of methods. The internal services 240 can include cryptography components, document securer processes, and an access control evaluation and creation infrastructure.

[0054] The methods that the server 200 exports to clients may depend on additional services with implementations that are dependent on a backend infrastructure of an enterprise system environment. The external service providers 260 can define a set of service provider interfaces that specify the connection(s) between the methods 250 and their execution environment. Upon initialization, the server 200 can load and initialize the set of service providers that are needed for this environment. The external service providers 260 can include default implementations and can be added to over time with additional implementations, tailored to different backend infrastructures, using the included service provider interfaces.

[0055] Example service providers are discussed below, but additional or alternative service providers are also possible. The definitions of the service providers are given in terms of interfaces that the service providers implement. These interfaces can be defined generically so that they can be implemented across a wide variety of systems. Thus, information that crosses system boundaries can be defined in simple terms to provide greater flexibility in implementation on various systems.

[0056] An authentication service provider can be used to authenticate a user. In the context of computer security, authentication is the procedure by which a program-machine confirms the identity of another machine, and/or the other machine's current user, from which a communication has been received. There are many types of systems in which authentication can be used, and there are many types of events that can trigger an authentication process, depending on the needs of a particular implementation. The authentication systems and techniques described herein can be used in a document control system as described, or in other systems.

[0057] FIG. 3 is a block diagram illustrating workflow in an authentication system. A client 310 can be communicatively coupled with a broker server 320 via a network 300. When the client 310 needs to take an action that depends on having an authenticated user, the client

310 can send a request 350 to the broker server 320. For example, when the client 310 needs to take an action with respect to a document 305, the client 310 can send the request 350. The request 350 can indicate to the server 320 that an update concerning the currently approved authentication process, for use in connection with the action, is expected by the client 310. The request 350 can include information indicating the action and/or one or more authentication processes already resident in a location local to the client 310; and the server 320 can determine, based on this received information, whether to respond to the client's request by sending an authentication process for use by the client 310.

[0058] Additionally, the request 350 can represent multiple communications between the client 310 and the server 320. The client 310 can first communicate to the server 320 that the action has been requested, and the client requests to know whether authentication is to be performed, and if so, how authentication is to be performed. The information identifying the server 320 and the document 305 can be included in the document itself, and the server 320 can determine whether user authentication is needed based on the information identifying the document 305 and the nature of the requested action. The server 320 can respond as to whether authentication is needed, and if so, the type of authentication to be used, including potentially multiple types of acceptable authentication mechanisms, from which the client 310 can choose which one to use. If the client 310 does not already have the specified authentication functionality, the client 310 can then request a corresponding authentication update.

[0059] The server 320 can be a dedicated authentication broker server, or the server 320 can provide other resources as well. For example, the server 320 can be a document control server as described herein, and various client-initiated operations (e.g., document viewing, revoking and securing) can effectively also be server-based operations in that completion of these operations may require contacting the server; such server-based operations initiated by a client can also trigger authentication using a dynamically delivered authentication process.

[0060] The server 320 can respond to the request 350 by obtaining an authentication process 315 and sending the authentication process 315 to the client 310. The authentication process 315 can be stored by the server 320 or by another server (e.g., a server in an enterprise system). Thus, authentication components can reside at the client 310, on the server 320, and optionally on a separate authentication server. Authentication can be handled via a service provider interface that allows the server 320 to be configured to use an existing enterprise authentication mechanism (e.g., password-based authentication), or even to implement a custom authentication mechanism that may be developed later (e.g. a biometric authentication, or a new smart card system). The authentication service provider interface can define

the methods that the server 320 uses to authenticate a user, and authentication service providers can be implemented for Windows and LDAP (Lightweight Directory Access Protocol) authentication, and also for one or more document management systems, such as authentication using the Documentum® Login Manager in the Documentum® content management system provided by Documentum, Inc. of Pleasanton, California.

[0061] The authentication process 315 represents a software program having instructions operable to cause a machine to perform operations effecting an authentication procedure. The authentication process 315 can become a component of the client 310 upon receipt or stand alone and communicate with the client 310. The authentication process 315 can be a plug-in to a document viewing application, such as the ADOBE ACROBAT® software provided by Adobe Systems Incorporated of San Jose, California. The authentication process can use an existing interface provided by the client 310 to communicate authentication information to the server 320 (e.g., the document viewing application can include a security handler component 317 that communicates with the authentication process 315, such as described further below). The authentication process 315 can be a client authentication library (e.g., a dynamic link library (DLL)) or a server service provider.

[0062] Thus, the client 310 can be transparently updated with a new authentication process as a result of sending the request 350 to the server 320. The specific mechanism(s) of authentication is therefore configurable, and end-to-end delivery of authentication components can be performed without the user being aware of the update. If an administrator changes the authentication procedure to be used for a document, all clients that attempt to perform an action that requires the specified authentication with respect to that document can be automatically and transparently updated to be able to authenticate using the newly specified mechanism. An authentication procedure can even be changed between sequential actions on a document, and thus a new request 350 can result in a new authentication process 315 being delivered for the same action to be performed on an already delivered document.

[0063] The authentication process 315 can implement an authentication procedure at the location of the client 310, interfacing and controlling any local hardware as needed (e.g., a biometric authentication procedure using biometric reading device), and the authentication process 315 can use an interface provided by the client 310 to communicate authentication information back to the server 320. The authentication process 315 can implement a wide variety of different authentication procedures, including multi-level and/or multi-factor authentications depending on the action being attempted. Because the authentication process 315 can be dynamically delivered in response to each request, an organization can readily change authentication procedures, adding new security features to a document control sys-

tem as needed.

[0064] The authentication process 315 can query a user at the client 310 for input (e.g., text, biometric data, etc.), encode the received input, and return the encoded input to an authentication provider on the server 320 (e.g., send the encoded input to the security handler 317 in the client 310, which forwards the information to the server 320). The server 320 can then handle authentication, either directly or in conjunction with an authentication server 330. In this pass-through authentication mechanism, the client 310 can provide credentials to the server 320, and the server 320 can work with a third party authentication system, such as LDAP or RADIUS, to authenticate the user. If authentication is successful, the authentication service provider can return an authenticated username.

[0065] Additionally, the server 320 need not be able to directly interpret client authentication information. Instead of the client 310 giving credentials directly to the server 320, the client 310 can first authenticate and then provide some resulting information to the server 320 to allow the server 320 to re-verify that the client 310 previously authenticated. For example, the authentication process 315 can contact the authentication server 330 to authenticate the user directly, and a receipt of authentication can be returned to the server 320. The server 320 can pass the receipt to the authentication server 330 and verify that there was in fact a successful authentication. Thus, the client 310 can provide credentials to a separate authentication system directly and then provide an authenticated token to the server 320, which can be used to verify the user's identity with the separate authentication system.

[0066] The server 320 can use multiple authentication service providers. The server 320 can dynamically deliver one or more authentication processes 315 to the client 310, as needed, using the interface described below. Such authentication process(es) 315 can be delivered securely to the client 310 and spoofing can be prevented, such as described below in connection with secure code library loading. The client 310 can also have one or more default authentication processes already available, such as an authentication library that can capture username-password text entry. Such default authentication process(es) can include support for user interface (UI) customization and a standard format for extracting this information within authentication service providers. Moreover, the client 310 can retain credentials for a period of time so that a user need not logon each time they perform an operation. An example of such retaining of client credentials to support offline access is described further below in connection with FIGS. 11-14.

[0067] Secure code library loading can be implemented to all the server(s) 320 to push one or more authentication libraries (e.g., DLLs, java bytecode, javascript, etc.) to clients to provide updates or customize clients without requiring any action (or knowledge) on the part

of the user while also preventing these authentication libraries from being spoofed on the client (e.g., by a Trojan horse program). A mechanism can be provided to verify the authenticity of the authentication libraries downloaded from the server 320. When the server 320 pushes an authentication library to the client, the server 320 can compute a hash of the library and also send this hash to the client 310, and/or the server 320 can sign the authentication library before sending it to the client. The hash can be retained locally at the client, and the client 310 can ensure the authentication library is valid by computing a hash of the authentication library and verifying it against the retained value at load time. Additionally, a selected set of libraries can be signed by the provider, or all the libraries can be signed by the provider, and the provider's public key can be retained at the client 310 (e.g., a DLL can be signed by Adobe when the client 310 is the ADOBE ACROBAT® software with the Adobe public key included).

[0068] FIG. 4 is a flow chart illustrating an authentication technique employed by a server. A request to take an action with respect to a document is received at 400. In response to the request, an authentication process is obtained at 410. The authentication process is sent to the client, at 420, for use in identifying a current user and controlling the action with respect to the electronic document based on the current user and document-permissions information associated with the electronic document. Thus, the authentication mechanism can be specified on the server and the appropriate code can be downloaded to the client dynamically, as needed, in a manner that is transparent to the client.

[0069] An authentication interface can provide either a text-based username-password description or a single authentication library. This can be implemented using two types of methods for authentication. The first method can take an opaque token (e.g., an uninterpreted byte string) as well as a username, although the implementation can choose to ignore either. The second method can take a username, password and optionally a third argument, which can specify the "domain", or a "connect string" if desired. The authentication provider can implement its own defense against brute force attacks, and can have the option to deny authentication even if the correct credentials are presented.

[0070] Implementations can also return an authentication reply that specifies whether the user successfully authenticated (verified). If verified is false, an additional error message indicating why it was not verified (e.g., no such user) can be returned; this error message need not be returned to the client, but can just be logged on the server (so as not to provide the client with helpful information that could be used to crack the authentication system). A token to be used in future authentication attempts can also be returned, although the server can ignore this. The username should also be returned for verified attempts such that the server can understand who has authenticated. The access control list (ACL)

service provider should be able to take this username and canonicalize it. The canonical form of the username can be consistently used across workflows, and the definition(s) governing canonical form(s) in the system can vary with implementation.

[0071] Because the client can authenticate via multiple methods, the server should be able to describe how the client should attempt to authenticate by default, or if authentication failed what method to attempt next. The authentication service provider can describe how authentication should occur - e.g., via a specific code library or via a basic text entry dialog being displayed to the user. If a code library is to be used, the server can communicate metadata about the code library to the client (e.g., a DLL's name, size, etc.). If a basic text entry dialog is used, the server can specify what the UI should look like to the user -- e.g., the title should say "Please enter your company LDAP password", and that only two fields, "username", and "password" are required.

[0072] In addition to the authentication systems and techniques described, document control systems and techniques can be provided. These can be combined with the described authentication or used separately.

[0073] FIG. 5 is a block diagram illustrating workflow in a document control system. A client 510 can be communicatively coupled with a permissions-broker server 520 via a network 500. A document source 530 can also be communicatively coupled with the permissions-broker server 520 via the network 500. The document source 530 can be a document repository (e.g., a document management system or a file system) and/or a document handling system (e.g., an email system). In general, the document source 530 can be considered one of two types: (1) a document source where a document 540 should be expected to be retained and accessible in the future, and (2) a document source where a document 540 should not be expected to be retained and accessible in the future (although it may be in practice).

[0074] When the document source 530 is of the first type, document-permissions information 550 can be retained at the document source 530 and sent to the permissions-broker server 520 when needed. Thus, the document-permissions information 550 need not be retained at the permissions-broker server 520 (although such information can be retained at the server 520 in a permissions-definition format specified for the server 520). When the document source 530 is of the second type, the document-permissions information 550 can be generated at the document source 530, at the permissions-broker server 520, or at the client 510, when the document 540 is secured to create a secured document 545, and the document-permissions information 550 can be retained at the permissions-broker server 520. The document-permissions information 550 can be an ACL that defines the types of actions that are authorized for the document 540. Moreover, document-permissions information can specify access permissions at a

level of granularity smaller than the document itself (e.g., controlling access to specific page(s), paragraph(s) and/or word(s) in the document).

[0075] The secured document 545 can be encrypted using an encryption key generated by the permissions-broker server 520, and the secured document 545 can include information identifying the server 520 and the document 545 (e.g., a link to the server 520 and a document identifier that is unique within the context of the server 520). The secured document 545 can be delivered to the client 510 in any manner (e.g., email, download from a document repository, received on a compact disc, etc.), and the secured document 545 can be a copy of another secured document (e.g., an attachment to an email forwarded from another source).

[0076] When the client 510 needs to take an action with respect to the secured document 545, the client 510 can determine that the document 545 is secured, extract the information identifying the server 520 and the document 545, and send a request 515 to the server 520 corresponding to the action and including the document identifying information. In response to this request, the permissions-broker server 520 can translate the document-permissions information 550 into second document-permissions information 555. The second document-permissions information 555 can be sent to the client 510 to govern the action with respect to the document 545 at the client 510. The client 510 can be a document viewing application, such as the ADOBE ACROBAT® software provided by Adobe Systems Incorporated of San Jose, California, and the document 545 can be a PDF™ document.

[0077] FIG. 6 is a flow chart illustrating a document control technique employed by a permissions-broker server. A request from a client to take an action with respect to an electronic document is received at 600. In response to the request, first document-permissions information associated with the electronic document is identified at 610. The first document-permissions information can be in a first permissions-definition format. The identified first document-permissions information is translated into second document-permissions information in a second permissions-definition format at 620. The second document-permissions information is sent to the client to govern the action with respect to the electronic document at the client at 630.

[0078] Referring again to FIG. 5, the first document-permissions information 550 can be in a first permissions-definition format that includes at least one type of permission information that cannot be fully defined in the second permissions-definition format used in the second document-permissions information 555, and translating between the two sets of information 550, 555 can involve translating based upon additional information associated with the request 515. For example, the first information 550 can include time-dependent permission information that cannot be fully defined in the second information 555 because the permissions-definition for-

mat includes no notion of time. But this time-dependent permission information can be defined in the second document-permissions information 555 for the limited purposes of the current request by taking the time of the request into consideration. If the first document-permissions information 550, in conjunction with the time of the request 515, indicates that the requested action is authorized, then this can be represented in the second document-permissions information 555; and likewise, if the first document-permissions information 550, in conjunction with the time of the request 515, indicates that the requested action is not authorized, then this can be represented in the second document-permissions information 555. When a subsequent action is requested, the translation can be performed again based on the time of the subsequent request.

[0079] As another example, the first information 550 can include user-dependent permissions information that cannot be fully defined in the second document-permissions information 555 because the permissions-definition format includes no notion of users. This user-dependent permissions information can include both user and group-based document control information and can be defined in the second document-permissions information 555 for the limited purposes of the current request by taking into consideration user-identification information obtained via the client 510. This user-identification information can be obtained using the authentication systems and techniques described elsewhere herein. When a subsequent action is requested, the translation can be performed again based on newly obtained user-identification information. Moreover, the multiple requests received by the permissions-broker server 520 can cause the server 520 to store information 525 relating to the actions taken at the client 510 with respect to the document 545. These actions can be associated with the username, and also with a network address (e.g., an Internet Protocol (IP) address) associated with the client (both as reported by the client and as reported by the server). Requested actions can also be considered actions taken, and the stored information 525 can be used by the server 520 to generate an audit of stored actions-taken information associated with the document 545, as described further below. The stored information 525 can also include actions performed and/or requested at either the server 520 or the document source 530 (e.g., actions performed at the file system, document management system, etc.), and a generated audit can include this information as well.

[0080] FIG. 7 is a block diagram illustrating workflow in a document control system integrated with a document repository 700. A permissions-broker server 730 can be used to secure documents in the repository 700 in a batch mode (e.g., when the server 730 is first installed) and/or as a step in a content management system (CMS) workflow. A securing client 720 can retrieve a document 710 from the repository 700. A document identifier 715 can also be retrieved and passed to the

server 730. The document identifier 715 can be used internally by the sever 730 to control actions with respect to the content. If the repository 700 is a CMS, the document identifier 715 can be the document identifier used in the CMS 700, and if the repository 700 is a file system, the document identifier 715 can be the URL (Universal Resource Locator) of the document.

[0081] The server 730 can communicate with the repository 700 using the document identifier 715 to obtain document-permissions information 740 (e.g., an ACL from a CMS or file permissions information from a file system). The document-permissions information 740 can be specific to the document 710 or can define permissions for multiple documents (e.g., a policy maintained by a document management system, or a set of file permissions maintained by a file system). The obtained document-permissions information 740 can be used by the server 730 to generate an initial ACL for the document 710. A set of data 750 that can include the initial ACL, the document identifier 715, and a key generated by the server 730, can be sent back to the securing client 720. The client 720 can use the set of data 750 to create a secured document 760, which is an encrypted version of the document 710. This secured document 760 can include the initial ACL, the document identifier 715, and the key packaged as part of the document 760.

[0082] When a client attempts an action with respect to the secured document 760 (e.g., attempts to open the document 760 or any copies of this document), the document identifier 715 can be retrieved from the document, sent to the server 730 and used to obtain the current ACL for the document 760, where the current ACL reflects the current state of the document in the repository 700. Thus, actions taken with respect to the secured document can be controlled based on document-permissions information defining current permissions for a source document in the document repository 700. The source document can be the originally secured document 760, or in the case where secured documents are not sent back to the repository 700, the source document can be the original document 710. The server 730 need not store document-permissions information, as this information can be retrieved from the repository 700 and translated whenever access to the document 760 is requested, although the server 730 may store the document-permissions information for other purposes.

[0083] FIG. 8 is a block diagram illustrating workflow in a document control system integrated with an email client 800. The email client 800 can be a plug-in to an email system and can be used to secure an attachment 810 to an email. When a user chooses to secure an email attachment 810, the email client 800 can prompt the user for the rules they wish to apply to the attachment and/or the rules can be generated automatically based on a recipient(s) list for the email. The rules can be converted into an ACL 830 at a securing client 820 and sent to a permissions-broker server 840. The server 840 can store the ACL and return a set of data 850, such

as described above. This data 850 can be used to create a secure attachment 860 that includes a document identifier, which may be generated and stored at the server 840, an initial ACL and an encryption key.

[0084] When a client attempts an action with respect to the secured document 860 (e.g., attempts to open the document 860 or any copies of this document), the document identifier can be retrieved from the document, sent to the server 840 and used to obtain the current ACL for the document 860, where the current ACL reflects the current state of the document ACL stored in the server 840. The sender of the email can interact with the server 840 to change the current ACL for the document 860, even after the email has been sent. Thus, actions taken with respect to a secured document can be controlled, and nature of the security on the document can be modified, even after the secured document has been distributed.

[0085] FIGS. 5-8 illustrate access control infrastructure as can be implemented in a document control system. In the context of the server described in connection with FIG. 2, an access control service provider can be implemented, where access control can be defined in terms of access control lists (ACLs). ACLs can map permissions (e.g., can print, can view, etc.) to principals (e.g., users and groups), and visa versa. The access control service provider interface can define the methods used by the server to map these principals into a canonical form that can be consistently used across workflows. Access control service providers can be implemented for various systems, such as NIS (Network Information Service), LDAP, and an email system (e.g., Majordomo, which is a public software program primarily running on UNIX machines to handle internet mailing lists). Moreover, the access control infrastructure can support shared ACLs (e.g., one ACL to be shared amongst multiple documents; such shared ACLs can be referred to as policies).

[0086] FIG. 9 is a block diagram illustrating a document control server 900 corresponding to the example of FIG. 2. The server 900 can support a variety of basic features, including: (1) Access Control - the ability to control who can access a document and what permissions they have; (2) Revocation - the ability to revoke a document so that it can no longer be viewed; (3) Expiration and/or validity intervals - the ability to specify time before which and after which the document cannot be viewed; (4) Document Shredding - the ability to make a document unrecoverable with respect to the document control server upon the document's expiration by destroying the document decryption key; (5) Auditing - the ability to audit actions performed on a document (e.g., viewing, attempted viewing, etc); and (6) Offline Access - the ability to access a document when offline. In addition, features can be easily added without changing the architecture.

[0087] An authentication service provider 910 can be implemented as described elsewhere herein, and an ac-

cess control service provider 930 can effect the access control infrastructure described. ACLs can include a set of Access Control Entries (ACEs) and a set of properties. ACL properties can apply to the ACL as a whole (e.g., expiration date). An ACE can map principals to rules and can include a list of principals, a rule, and a validity period for the ACE. When an ACL is evaluated, only ACEs that are within their validity period need be considered. Validity periods can allow different users and groups to be granted permission to view a document at different times. For example, an ACE can specify that "only members of the public relations staff may view a document before its release date, after which anyone can view the document."

[0088] Rules can include of a set of properties and granted and denied permissions. These permissions can be specific to a viewing client application (e.g., the ADOBE ACROBAT® software) and/or server defined. Additionally, permissions, like properties can be extensible, so new ones can be added without changing the ACL format.

[0089] The server 900 can have its own simple mechanism that allows users to specify Access Control Lists using a Securing Client interface without the use of any external ACL mechanism. Additionally, third party ACL/rights specifications can be translated to the internal ACL format used by the server 900. The server 900 can integrate with other systems' access control facilities (e.g., Document Management Systems, Database Systems, File Systems, etc), leveraging the functionality in these systems.

[0090] The server 900 can support integrating with diverse user and group repositories that may contain incomplete information, and the server 900 can be enabled to efficiently access this information in a canonical user-centric manner. Facilities for manipulating ACLs on both the server 900 and a client 980 can be provided. The server 900 can verify ACLs to ensure they are valid before a document is secured, either using a server-based document securer 960 or a client-based document securer 990. ACLs can be extensible and can allow opaque third party permissions. Moreover, securing of documents can be done in an online-fashion, connected to the server 900, because the server can verify ACLs.

[0091] The server 900 can associate ACLs with documents in order to specify which principals (e.g., users and groups) have which permissions for a document. A principal can have multiple names; however, a principal should also have a distinguished canonical name. One of the tasks of the server 900 can be translating the various names of a principal into its canonical name. While both permissions and properties can describe authorized operations, permissions can be boolean valued and properties can be of a variety of types. Permissions can be granted if explicitly granted and not explicitly denied; undeclared permissions can be implicitly denied.

[0092] Each document can be associated with a sin-

gle ACL. Typically this relationship can be 1:1, but in the case of policies this relationship can be N:1, where multiple documents share the same ACL. The electronic document file can contain an immutable snapshot of the ACL dating to the time of securing. The server 900 can also maintain a copy of the latest ACL, which can be modified by authorized individuals. The server 900 can canonicalize ACLs (e.g., translate all principal names to their canonical forms) before they are used. This can be done whenever ACLs are created or modified (e.g., at the time of securing, or when ACL definitions are changed). Once ACLs are in canonical form, it can be much simpler to evaluate ACLs on both the clients 980 and the server 900 since determining membership within groups as well as determining relevant authorizations for specific authenticated users can be done via basic string matching.

[0093] The server-side evaluation of ACLs for a specific user at a specific point in time (e.g., for online viewing, revocation, document audit retrieval, etc.) can be implemented within the server 900 directly. The server 900 can examine the ACL, looking for ACEs that are currently valid and that also contain either the authenticated user or a group in which s/he is a member, and then extract the permissions and properties. The server infrastructure to handle canonicalization within the server 900 can have three tiers. A first tier can be an in-memory cache in the server 900 that maps non-canonical principals into their canonical forms. A secondary persistent cache can store canonical mappings and user-in-group information; this cache can potentially be used across multiple servers 900. The third tier can be the access control service provider 930.

[0094] The access control service provider 930 can include a set of principal modules that provide the canonical form of some set of non-canonical strings. These principal modules can also specify whether the canonical form corresponds to a canonical group or a canonical user. However, the architecture need not assume that a specific principal module will generally know all answers, or be able to give a complete answer about a specific non-canonical string. To support multiple domains of expertise within the context of user and group repositories, each principal module can publish the domain(s) over which it is the authority. The process of canonicalization, which can be implemented within the server 900 directly, can take a non-canonical form and iteratively refine it by querying modules with authority until one declares the returned value as canonical.

[0095] Methods 970 in the server 900 can be authenticated-user-centric, because a typical scenario involves the server 900 determining whether a specific user has permission to perform an operation, taking into account what groups s/he might be in. Many third party group mechanisms organize group membership accessible by "who are members of a group?", but not "which groups contain a specific user?" Moreover, in many cases groups may contain non-canonical forms of users.

Thus, the output of group repositories may not be directly usable by the server 900, and a translation intermediary can be employed.

[0096] A very low common denominator can be assumed for group providers. A group provider can be expected to be able to provide a list of known canonical groups. Thus, valid groups can be those in the union of known groups specified by group modules. Group modules can also provide membership information organized in a group-centric manner, which can be an efficient approach given the implementation of many existing repositories.

[0097] The server 900 can have the capability to batch preprocess group information for subsequent use within the system. For example, one server in a group of servers can run such a batch operation on a daily basis. This can be implemented in the server core and can involve enumerating all groups, canonicalizing members, examining group nesting and computing the transitive closure. Most of the transitive closure computation can be within a storage provider 920, since it is natural to perform these types of operations using database systems.

[0098] A principal can be either a user or a group. Principals can be represented as strings. Groups can contain principals. Principals can have many alias expressions that can be evaluated and reduced to a primary canonical form. Users and groups can be of multiple domains. A convention involving the name@sub.domain.com format used in email addresses can be adopted, even if the document control system integration is not email-based. Moreover, the specification of what the canonical form should be can be left undefined in the general system, as this specification can be integration-dependent. Examples in a particular integration context can be as follows: "herbach@company.com" is the canonical form for many strings, including "jonathan_herbach@corp.company.com" and "jherbach@company.com"; likewise, "atg@company.com" is the canonical form for "atg@sea.company.com".

[0099] An access control service provider interface can include principal providers, which can be divided into two subtypes: user modules and group modules. The goal of these modules can be to provide canonical information and group membership information. A principal provider can translate a principal, to the best of its ability, into canonical form. The principal provider can indicate whether the returned value is in canonical form, whether it is known to be a group or a user, and how long the returned result can be considered valid in a cache. A principal provider can have a domain of authority, specified as a set of regular expression definitions, and a group provider can enumerate all the groups it knows about in its domain of authority.

[0100] To support the various server methods 970, user and group information can be provided logically, as there might be multiple sources of such information. Thus, there can be several User Modules and several Group Modules. From a high level, each one can be con-

figured differently, can interface with different backend systems, and can be an authority over possibly multiple domains. Moreover, defining different modules as domain authorities can assist in providing extranet support.

[0101] Configuration of the principal modules can describe the appropriate class file. Each module can also have some module-dependent configuration information, such as connect strings and preferences, as well as infrastructure to configure what the authorities are. Different implementations can also have a rule governing pre-processing and post-processing to facilitate integration with the rest of the system.

[0102] An ACL manager 940 can contain code relevant to loading an arbitrary number of principal providers. FIG. 10 is a block diagram illustrating example details of the server from FIG. 9. The server can have a primary in-memory cache, handled by an ACL manager 1010, for group membership or canonical mappings. The user can store within memory the recent canonical mappings such that the service providers need not be called for common requests.

[0103] The ACL manager 1010 can also include cross-method code, and an ACL Service Provider Manager 1020 can be a transparent interface to storage-level (e.g., cross-server) caching. Queries to the ACL Service Provider Manager 1020 can first result in checking whether a storage provider 1030 has the necessary information, and return that. If not, the ACL Service Provider Manager 1020 can issue queries to user and group modules 1040 and attempt to persist as much information to the storage layer as possible. Cache entries can be cleaned as per an expiration associated with the canonical result returned (e.g., as specified by either the storage provider or the principal modules).

[0104] Referring again to FIG. 9, a storage service provider 920 can provide an interface that describes a collection of methods that the server 900 uses to create and retrieve data in persistent storage. This interface can be the largest service provider interface in the system and can grow further as new integrations and features are implemented in a document control system. The storage service provider 920 can provide methods in the following areas: (1) Allocation of document tickets - each document that is secured on the server can be given a ticket with a GUID (global unique identifier); (2) Recording document revocation; (3) Saving encryption keys for users, groups, documents, and the root server keys; (4) Caching user alias and group membership data; (5) Auditing user access and securing; (6) Management and storage of named ACLs or policies; (7) Storage and retrieval of the current ACLs for documents; (8) Creation of initial ACLs for documents.

[0105] The storage provider interface can be designed to allow multiple implementations across a wide variety of backend systems. This can be done using a generic relational database implementation, which can work with both ODBC and JDBC. In addition, the storage

provider interface can be designed to support an implementation for a content management system, such as the Documentum® system. Ticket generation can be straightforward. For example, this can be implemented by having an integer in the database that is incremented on each reservation. Document revocation can be defined as the ability to revoke a document based upon its ticket and to separately query whether the document associated with a given ticket has been revoked. The storage provider can also store and retrieve keys, which can be arbitrary byte arrays, by name.

[0106] The storage provider can also provide storage for user alias and group membership data. Alias and membership information can be used to evaluate access control lists; the storage provider 920 can be used as a cache to help ensure reasonable performance even if the access control service provider 930 is not capable of providing efficient access to this information. For example, in the limiting case, the access control information might come from flat files that provide the required data. When caching user and group alias information, the storage provider can perform retrieval queries based upon a principal, much like user and group providers. The data returned should be of the same format, also providing an indication of the validity. The goal can be such that when the server uses user alias or group membership data, the server should not distinguish whether the data provided is real-time or a cached version.

[0107] For a given user or group, the canonical name of the user or group can be obtained. For a user, all of the groups to which this user belongs can be obtained. Changes to alias data can be immediately visible. Changes to the group membership cache may be more complicated, because of transitive closures computation (group memberships of groups that contain groups). Because of this, group content changes may not be immediately visible if the server is currently computing the transitive closure of groups.

[0108] Document securing operations and document access attempts (whether successful or not) can be audited through auditing methods of the storage provider 920. In addition to defining the methods to record securing and access events, the interface can also define a couple of query methods on the audit history - querying by document ticket and by user. The storage provider can also implement methods that allow ACL creation and modification. These methods can be used to keep auditing history information. Multiple implementations of the storage service provider 920 can be implemented as needed, including using a relational database and/or using existing document management system notions of audit logs (e.g., Documentum® audit trail objects).

[0109] The storage provider 920 can store and retrieve ACLs by name. An ACL can be a private ACL (e.g., for a particular user) or a public ACL. Public ACLs represent policies that are intended to be shared across multiple documents secured by various users. The stored representation of an ACL can be a matter of con-

cern only to the storage provider, as the provider implementation can be designed to simply take ACLs as arguments and return ACLs as results; the ACLs can be described in terms of an AccessControlList interface.

[0110] The storage provider can have a set of methods to create, update, delete, and retrieve ACLs. The methods can take arguments describing either a named ACL or a policy (e.g., a public ACL). There can also be methods to associate a stored ACL with a given document (via the ticket GUID). When associating a given document with an ACL, ticket data can also be stored. This ticket data can be specific to a particular document and can be used to store document-specific information like the date when the document was secured as well as which principal secured the document. An ACL shared amongst documents can also specify controls relative to the time of securing or to the person who secured the document. The ticket data can also be used by the securing client to provide information corresponding to the service provider. For example, in a Documentum® system integration the ticket data can provide the Documentum® GUID for the source document. The service provider information can also be a byte sequence received from the service provider including a set of name/value pairs that capture appropriate informational aspects of the document corresponding to the service provider.

[0111] In addition to the ability to retrieve ACLs by their name, the server can also retrieve an ACL for a specific document. When retrieving an ACL for use, the server can optionally provide a principal as a parameter. This provides a hint, allowing an optimized storage provider to return the subset of an ACL that is relevant for that particular principal.

[0112] When creating and storing an ACL, there is also the opportunity to pass through service-provider specific data that was presented to the securing client. This can provide an end-to-end mechanism to give a hint to the service provider on what specific ACL this document refers to. This is analogous to the capability described above in connection with the ticket data, but may be specific to an ACL as opposed to a document.

[0113] The storage providers need not interpret ACLs. The storage provider can simply store and retrieve ACLs without doing any interpretation of them. When a document is created it can be given an initial ACL, which can be stored in the document and used for offline access control if no other ACL for the document exists locally at the client. The storage interface can provide the methods by which these current and initial ACLs are passed back to the securing or viewing components of the server. In general, there can be two main cases: (1) the content being secured does not have any separate identity outside of the document control system (e.g., the content is an email attachment); (2) the content does have an identity outside of the document control system (e.g., the content is a PDF™ rendition of a document inside a Documentum® repository). In this latter

case, the service provider should be able to dynamically control access to the content in terms of the current rules the repository applies to the object from which the content was derived. Moreover, once an ACL has been saved, it can be modified by the owner, or by a system administrator in the case of a policy.

[0114] Both the initial and the current ACL can be generated by the storage service provider, and access control for the content can be mediated in terms of the access control on the underlying object. Otherwise, the management of the content may be precisely the same, in both the online and offline case. In addition, a Boolean supportsProvider method can be provided that the client can use to see what service(s) are supported by the service provider. The client can thus have an expectation of which service provider it can use, and can determine from the supportsProvider method if this service is actually supported by this document control server configuration (e.g., this determines what set of name/value pairs can be legally included in the service provider information in the ticket data). If supportsProvider() is true for some service, then the remainder of the interface should be implemented. Thus, a customer could use the same server both to protect content in a document repository and to protect email attachments.

[0115] The server 900 can also include a cryptography component 950, which can have duplicate implementations that take advantage of various native cryptography components (e.g., Java Cryptography Extension or .Net Crypto components). In general, a document control server uses several cryptographic primitives. These cryptographic primitives' implementations can be placed behind general interfaces, allowing the implementations to be changed (e.g., change key sizes, etc.) as needed, such as to add security features and/or to address the needs of specific enterprises. Additionally, these cryptographic primitives' implementations can use standard cryptographic operations as well as custom operations.

[0116] The interface of the cryptography component 950 can provide support for the following primitives: (1) symmetric encryption and decryption (e.g., 128-bit AES (Advanced Encryption Standard) and/or 128-bit RC4 (Rivest Cipher 4)); (2) public key encryption and decryption plus signing and verification (e.g., 1024-bit RSA); (3) message authentication code (MAC) used to provide document integrity (e.g., the one-way HMACSHA1 hash function with a 128-bit key); (4) a secure hash function for which it is computationally infeasible to find two messages that hash to the same value (e.g., SHA1); and (5) random number generation used to create cryptographic keys and introduce randomness into messages (e.g., the Secure Random number generator provided with the .Net framework for a .Net implementation and the java.SecureRandom class for generating random numbers in a Java implementation). These cryptography primitives can be implemented in Java using the Java Cryptography Extension (JCE) mechanism and in one

of the .NET languages using the .Net Service Provider mechanism. This cryptography interface and the cryptography implementations should also be used on the clients, as both the clients and the servers in the document control system can secure and access documents using these cryptography techniques. The cryptography interface can also be implemented in C++ for any cryptographic operations used on clients written in C++.

[0117] FIG. 11 is a block diagram illustrating an offline document access model as can be used in a document control system. A client 1110 can be communicatively coupled with a document control server 1120 via a network 1100. The document control server 1120 can provide multiple offline usage models, including a lease model similar to traditional offline access models, where the user must be online the first time a document is accessed and can subsequently access the document offline for a specified period of time, i.e., the lease period. In addition, the document control server 1120 can provide an initial access model, where the user can be offline when the document is accessed for the first time. As used herein, the term "online" means the client 1110 can communicate with the server 1120; thus, the client 1110 is connected with the network 1100, and the server 1120 is operational, when the client 1110 is online.

[0118] In general, the client 1110 and the document control server 1120 periodically synchronize to update any changes to offline access information retained at the client 1110, where this offline access information can effectively pre-authorize the client to allow actions with respect to secured documents that have yet to be accessed while the client 1110 is connected to the network 1100 (e.g., a secured document received via email at the client but not yet opened). The client 1110 can send a request 1130 to the document control server 1120. The request 1130 can be for an update to its offline access information. For example, an agent can be provided with the client 1110 that periodically connects to the server 1120 and downloads offline access information; this synchronization operation can happen silently in the background without a user of the client 1110 being aware of the updates; the next time the user attempts to open a document, the downloaded offline access information can be used by the client for future access while offline.

[0119] The request 1130 can be any type of request sent to the server 1120 periodically, such as a request from the client 1110 to take an action with respect to a document 1135, which may be located at the client 1110 or elsewhere and may be a secured document or not. The server 1120 can verify an authenticated user at the client 1110 in connection with the request 1130, and this verification of an authorized user can cause the synchronization operation to initiate. For example, the server 1120 can be a server such as any described above, and the synchronization operation can piggyback on other operations that use authentication (e.g., when a user attempts to access or secure a document while on-

line). Alternatively, synchronization can occur without prior authentication; the server 1120 can encrypt the offline access information using the user's public key so that only the user can decrypt them; the encrypted offline access information can be retained by the client 1110, and when the user next attempts to open a document, the retained information can be decrypted and used to update the client's secure local database as described further below.

[0120] When the client 1110 synchronizes with the server 1120, the server 1120 can send offline access information 1140, which includes a key 1145 associated with a group of users to which the current user belongs (a picture of a key is used symbolically in the figures to represent one or more encryption keys). The key 1145 can be used to access a secured electronic document 1150 while offline by decrypting a second key 1155 in the electronic document 1150. The electronic document 1150 can include content encrypted with the key 1155, and the electronic document 1150 can include the key 1155 encrypted with the key 1145. Alternatively, there can be one or more levels of indirection in this key encryption relationship. For example, the key 1145 can be used to decrypt the key 1155, which can be used to decrypt another key that is then used to decrypt the content of the document 1150. Regardless of the number of levels of indirection and the number of keys employed, the key 1145, which is associated with a group of users, can be used to access the secured electronic document 1150 while offline by decrypting a second key 1155 in the electronic document 1150. Additionally, the offline access information 1140 can include other group-specific keys, one or more user-specific keys, at least one set of document-permissions information associated with multiple documents (e.g., a policy as described above), and a document revocation list.

[0121] The synchronization operation can also involve the client 1110 sending back to the server 1120 an offline audit log 1160 of operations performed by the client while offline. Thus, the client can periodically synchronize with the server to upload audit log messages that have been retained locally and to download the latest revocation list and any updates to policies. In a system employing ACLs as described above, all new ACLs need not be downloaded with each synchronization because of the potentially large number of ACLs in the system. The document control system can provide a constrained set of guarantees as to the freshness of data. The guarantees used can be as follows: (1) Each document-specific ACL and policy specifies a period of offline validity (e.g., a number hours or days for which the document-specific ACL is valid before another synchronization with the server is needed, and after which, the document may not be viewed offline without synchronization). (2) At each synchronization, all revocations and policy updates are synchronized with the client. Thus, a policy or revocation list can be at most a specified number of time units out of date with respect to a par-

ticular document. Moreover, the synchronization can also send a current ACL for any document being accessed while online.

[0122] FIG. 12 is a flow chart illustrating a synchronization operation as performed by a server. A request is received at 1200. In response to the request, the server determines if an update is needed at 1210. For example, the server can compare a time of last recorded client-synchronization with a time of last change in user-group information for the user, or the server can compare current user-group information for the user with received user-group information for the user from the client (e.g., the client can identify to the server its currently retained user and group keys, and the server can respond based on whether any changes to the client's retained keys are needed).

[0123] If an update is needed, the server sends offline access information at 1220. This can involve the server sending the client a list of the keys to remove and the keys to add locally. If no update is needed, the server sends a validation of the current user-group information at 1230. This indicates to the client that current offline access information is valid, and the client and server are synchronized as of the current time. Additionally, when the server sends the offline access information at 1220 or revalidates the client's offline access information at 1230, the server can also send a server-reference time to be recorded at the client and used in determining when a client-server synchronization is needed again in the future. Finally, the server receives an offline audit log from the client at 1240. Thus, the server can generate audits, as described above, that include information relating to actions taken with documents while offline.

[0124] FIG. 13 is a flow chart illustrating a synchronization operation as performed by a client. Offline access information, including a first key, is received, and an offline audit log is uploaded to a server when the client is connected to the network at 1300. The client retains the offline access information at 1310. Cryptographic keys and other sensitive information can be retained locally on the user's machine in a secure manner, such that an attacker can not gain easy access to such information.

[0125] Security may be provided by encrypting the files with a cryptographic key stored in tamper-resistant hardware, such as a smartcard or an embedded security chip, such as those that ship with some laptops provided by International Business Machines Corporation of Armonk, New York. If hardware tamper-resistant storage is not available, software obfuscation techniques may be used to provide some security. The data retained at the client can include user and group private keys, a document revocation list, updated ACLs for policies, updated ACLs and security data for documents the client has accessed while online, and an offline audit log of operations performed by the client while offline.

[0126] A request to access a document is received when the client is not connected to the network at 1330. A check is made to determine if a recent server synchro-

nization has occurred at decision 1340. For example, the client can check whether a difference between a current time and a receipt time of the offline access information exceeds a server-synchronization-frequency parameter. The server-synchronization-frequency parameter can be specific to the document to be accessed. Moreover, determining the current time can involve comparisons between the last known synchronization time and the local system clock.

[0127] If a synchronization with the server has not occurred recently enough, the client prevents access to the document at 1350. If a synchronization has occurred recently enough, the first key is used to decrypt a second key in the document at 1360. Actions with respect to the electronic document can be governed based on document-permissions information associated with the electronic document at 1370. Governing actions with respect to the electronic document can involve obtaining the document-permissions information from the electronic document itself. Governing actions with respect to the electronic document can involve identifying a document policy reference in the electronic document, and obtaining the document-permissions information retained locally, based on the document policy reference. Additionally, an offline audit log, which can record both document access and attempted document access, can be maintained at 1380.

[0128] FIG. 14 is a block diagram illustrating components of a secured document 1400. Included within the secured document 1400 can be an encrypt dictionary 1405. The encrypt dictionary 1405 can include encrypted keys, which can be used to access the content of the document 1400, and an address (e.g., host name, port number, and connection protocol) of the server to contact when online. The encrypt dictionary 1405 can be embedded within the encrypted document 1400 in a location that is not encrypted by the document key used to encrypt the document (i.e., used to encrypt the document content).

[0129] An example encrypt dictionary 1410 includes document permissions information 1420 (e.g., the initial ACL described above) and one or more encrypted document keys 1430. The document key used to encrypt the content of the document 1400 can be encrypted multiple times using group keys and user keys, and these encrypted document keys 1430 can be included in the encrypt dictionary 1405 in the secured document 1400. A document control server can dynamically generate and maintain user and group keys for the user and groups in a document control system. By including the encrypted document keys 1430 and the document-permissions information 1420 in the document 1400, offline access can be supported by providing the appropriate user and group keys to the client using the synchronization operation described above.

[0130] Another example encrypt dictionary 1440 includes a document key 1450, an ACL 1460, a document ticket 1470, version information 1480 (e.g., a format ver-

sion string), and encrypted session keys 1490. The document key 1450 can be a random 128-bit key generated by the document control server and used to encrypt the document content (e.g., using RC4 or AES encryption). A portion of the encrypt dictionary 1440 can be encrypted using a generated session key, and a MAC can be used to detect any modification of the encrypt dictionary. The encrypted session keys 1490 can be the session key encrypted multiple times using the group keys and the user keys. Additionally, the session key can be encrypted with the server's public key.

[0131] When a user attempts to open a document offline, the client can check to see if the session key for the document has been encrypted with the user's key or the group key of any group of which the user is a member. The client can obtain the user's key and keys for all groups of which the user is a member during synchronization with the server. The appropriate key is then used to decrypt the information in the document's encrypt dictionary. The client can then evaluate the ACL in the same way ACLs are evaluated on the server to determine what permissions the user has. The client's revocation list can be checked, and if the document has not been revoked and has not expired, the document can be opened and the user's access to the document can be audited locally.

[0132] This initial access model allows a user to be offline the first time they access a document. When the document 1400 is secured, the initial ACL for the document can be embedded, immutable, in the document. When a user attempts to open the document, the embedded ACL can be used to determine whether they have access. The document 1400 can still be revoked or expire even though an initial ACL is kept within the document. Moreover, the current ACL for the document 1400 maintained elsewhere can be updated, and this ACL can be used when the client is online, as described above.

[0133] When a user accesses a document online, the current ACL, which can be stored on the server, can be retained on the client and used for that access. The retained ACL can then be used for future offline access to the document. When the client obtains the updated ACL from the server, the client can also obtain the document session key, separately encrypted with the key of each user and group that can access the document. Both the ACL and the encrypted keys can be secured in a manner similar to that initially embedded in the document.

[0134] Moreover, the document permissions information 1420, 1460 in the document can include a policy, i.e., a document policy reference or identifier. Thus, the client can identify a document policy reference in the electronic document while offline, and obtain the document-permissions information of the policy, retained locally, based on the document policy reference. As the document control system can guarantee that all policy updates are reflected on the client with each client-server synchronization, an administrator can change a pol-

icy and know that within a bounded amount of time, the change will be reflected on all clients that are still providing access to any documents.

[0135] In addition to the initial offline access model described above, a traditional lease model can also be used in the document control system to provide additional flexibility. In this model, the first time a user accesses a document from a particular machine, they must be online. At that time, they receive an offline lease, which allows them to view the document for a specified period of time offline before the lease must be renewed. Such a lease model can be implemented in the document control system described by embedding an initial ACL allowing access to no principals, and employing a validity_interval that specifies how long an ACL can be retained on the client before a new one needs to be fetched from the server. Additionally, the document control system can be configurable to enable a no-offline-access model in which the user must be online in order to access a document; in this case, the keys needed to open the document need not ever be retained on the client.

[0136] The document control system can provide all of the following security guarantees together as well, generally subject to the accuracy of client time. (1) Policy Modification - A policy modification is guaranteed to be reflected on each client within the offline_validity_interval specified in the policy since all policies are synchronized at every synchronization operation. (2) ACL Modification - A (non-policy) ACL that has been modified will be reflected on the client only if it is viewed while online. Retained non-policy ACLs are guaranteed to be dropped from the client within the validity_period if specified in the ACL. (3) Revocation - A document that has been revoked is guaranteed to be unviewable by all clients in the system within the offline_validity_interval, specified in the document's ACL since revocation is synchronized with the client at every synchronization operation. (4) Expiration - A document that has expired will be unviewable on the expiration date regardless of whether the user is online or offline. (5) Expiration modification - Expiration is specified in the ACL, and so expiration modifications are reflected as per-Policy or per-ACL modifications. (6) User or Group membership modification - If a user's key is revoked (e.g., because they leave the company) or if the user is removed from a group, it can be guaranteed that the user will not be able to view a document that they no longer have access to within the offline_validity_interval for the document.

[0137] FIG. 15 is a flow chart illustrating a document information delivery technique employed by a server. A request for a client to take an action with respect to a first electronic document is received at a server at 1500. In response to the request, information associated with the first electronic document is identified at 1510. The associated information can indicate a second electronic document that is different from and associated with the

first electronic document. This information can associate two or more documents and can describe the relationship(s) between them; this association information can be stored at the server, such as in a table or a database. Information concerning the second electronic document is related to the client at 1520 to facilitate the action to be taken.

[0138] Relating the second document information to the client can involve sending the second document information to the client to allow selection of one of the first and second documents with respect to the action. Relating the second document information to the client can involve obtaining the second electronic document, and sending the second electronic document to the client to allow taking of the action with respect to the second electronic document instead of with respect to the first electronic document. The second document can already exist or may need to be generated in whole or in part, which can be indicated by the associated information indicating the second document.

[0139] FIG. 16 is a block diagram illustrating workflow in a document control system. A client 1610 can be communicatively coupled with a document control server 1620 via a network 1600. The client 1610 can send a request 1630 to the document control server 1620, where the request 1630 relates to an action to be taken with respect to a document 1640. The server 1620 can check information 1645, which can be stored locally or elsewhere, that is associated with the document 1640 and indicates a second document 1650. The server 1620 can then send information 1655, which can be information about the second document 1650 and/or the document 1650 itself.

[0140] The client 1610 can force a user to view the second document 1650 based on the information 1655. For example, the second document 1650 can be a later version of the first document 1640, and the information 1655 can include document-permissions information specifying that the action is not permitted with respect to the first document 1640. The first document 1640 can be replaced with the second document 1650 (e.g., opened in place of the first document and/or written to storage over the first document) by the client 1610, including potentially without the knowledge of the user. The second document 1650 can also be a different language version (e.g., a French version of an English original) or a different format version (e.g., a different file compression and/or encryption scheme) of the first document 1640.

[0141] Obtaining the second electronic document 1650 at the server 1620 can involve generating at least a portion of the second electronic document 1650 (including potentially generating the entire document 1650), or the document 1650 can be a pre-existing document. The associated information 1645 can include user-based association information, and obtaining the document 1650 can involve obtaining the document 1650 based on the user-based association information

and an identified user at the client 1610. The document 1650 can be customized for a particular user, the user's location and/or the user's time of access (e.g., the document 1640 can be a stub document that is already identified as outdated when sent, and when this stub document is opened, each user can automatically receive a new document generated specifically for that user at the time of the access attempt, i.e., the stub document looks like and can be manipulated as a regular document in an operating system, but is always current when opened while online). Customization of the document 1650 can be done at the server 1620 or elsewhere. The user can be identified as described above, and the document control system can also employ the systems and techniques described throughout this patent application; the documents 1640, 1650 can be secured documents as described above.

[0142] FIG. 17 is a flow chart illustrating a document information receiving technique employed by a client. A locally retained distributed document is opened at 1700. The distributed document can be a secured document, as described above, that identifies a document control server to contact. A document control server identified from the distributed document is contacted at 1710. The server can determine whether the distributed document is the appropriate document, or if a different related document should be used instead. Use of a second document in place of the distributed document is forced at 1720, with respect to a document action, based on information received from the document control server.

[0143] A document control system can thus address both issues of document security and version management in one system. If a different version of a distributed document should be viewed in place of the distributed version, this can be defined and controlled in a document control server that also handles document security for distributed documents. An author of a document can specify that a distributed version of a document is outdated, and a newer version should be viewed instead. Moreover, an author can easily control multiple versions of a document and user-based definitions of who should view which version.

[0144] An author or administrator can designate which documents are appropriate versions for which recipients, including the possibility that two users receive entirely different documents with different content and which are different document versions in the sense that they both relate to an originally distributed document. Version relationships among documents can be specified using the document identifiers generated for document security purposes. The version relationships can be defined using a directed graph in which each node is a version, and the directed edges indicate which versions take precedence. Each edge can also indicate to which users it applies. A graphical user interface for displaying diagrams can be used to define the version relationships, such as by drag and drop operations to specify which versions become outdated in favor of oth-

er versions.

[0145] In the context of different sequential versions of a document, where each document can be revised and the system can ensure that each user only views the latest version of a document, the notion of revocation in the document control system can be extended to include whether a document has been replaced with another. Thus, upon opening a document, in addition to checking whether users have access to perform traditional actions on the document (e.g., print, etc.), a determination can be made as to whether the user should have access to a specific version of the document. The server 1620 can store information about where documents can be found, including potentially providing an additional repository service where documents that are being persistently versioned can be stored.

[0146] In the case where each user can view a different version, a similar approach can be used, with the addition of the ability to specify intersecting user/groups (e.g., "instead of version zero, all employees should see version A; all managers should see version B; and an executive should see version C", where additional version relationship information specifies that the executive can open the subordinate versions A and B in addition to version C). Rules for resolving conflicts can be provided.

[0147] The systems and techniques described herein can be combined in a comprehensive document control system employing multiple document control servers. Referring again to FIG. 9, the document control server 900 can implement the various techniques described, in combination. To increase system security, all client-server communications can be over Secure Socket Layer (SSL), which encrypts the communications and provides server authentication, and/or securing of documents can be done using client-side securing. The server 900 can be physically secured from an attacker and can sit behind at least one firewall. All sensitive state information in the server 900 can be encrypted before it is persisted to stable storage; the encryption key used for this can be embedded in the server code, hidden in obscure system resources and/or contained within a tamper-resistant cryptographic module. Moreover, on the client side, a user's logon credentials can be cached to avoid repeated authentications for multiple consecutive operations that require authentication. Cached credentials can be signed by a server private key, dedicated to this purpose, and reside on the client; the signed credential can include an expiration date to limit its validity period and can be presented when the client attempts to authenticate against the server 900.

[0148] As mentioned above, documents can be secured either at the server or at the client. A document can be converted from one format to another (e.g., from Microsoft Word to PDF™) before securing; the document control system can be integrated with a PDF™ creation service for this purpose. The securer component 960, 990 can be a wrapper around a PDF™ library that

takes a PDF™ document as input as well as an encryption key and a set of name/value pairs that represent information to be embedded in the PDF™ document's encrypt dictionary. The securer can encrypt the document with the provided encryption key and embed the specified information in the document. When the securing is performed on the server 900, the securing can be done in a separate process - a pool of such processes can be kept available so that multiple securing requests can be simultaneously satisfied, and the maximum number of such processes can be a configuration option for the server 900. These securing processes can be terminated after some number of successful securing operations, which number can also be a configuration option, or after any unsuccessful securing operation.

[0149] FIG. 18 is a block diagram illustrating document securing workflow in the document control server of FIG. 9. Securing a document can generally involve two high-level operations: preparing system state associated with securing of a document, and embedding relevant information into the document and encrypting it. Preparing state can be a joint operation between the securing client, specifying how a document should be secured, and the server, which can prepare the system for the secure document. Embedding information into the document and securing can be done either on the server (e.g., the unencrypted document is sent up to the server at time of securing and then the encrypted form is returned to the client), or on the client (e.g., the client has the components necessary to encrypt the document).

[0150] The securing client can prepare a specification of the desired security for the document to be secured. This can involve end-user interaction in a client, such as an email application like Outlook® software, provided by Microsoft Corporation of Redmond, Washington. The client can connect to the server via the RPC, authenticate, and send information up to the server (1800). If the system is using server-side securing, the client can send the unencrypted document and the securing specification up to the server. If the system is using client-side securing, then only the specification need be sent.

[0151] The server can authenticate the user, ensuring that he has permission to secure a document (1805). The service provider can provide a ticket (GUID) for the document (1810). The Access Control List specification can be given to the Access Control Manager so it can canonicalize the principals and possibly validate permissions (1815). The ACM can first attempt to use an in-memory cache of canonical mappings. The storage provider can be queried for other cached canonical mappings (1820). Principal providers can be queried for all non-cached noncanonical entries (1825). The canonicalized ACL can be persisted in the storage provider to allow for subsequent modification of the ACL (1830).

[0152] The information to be encrypted and stored in the document (e.g., ticket and ACL) can be provided to the Crypto Service Provider (1835), which can create a document key that will be used to encrypt the document.

If document shredding is not desired, then document key, ticket, and ACL can be encrypted using the server public key. If shredding is desired, then the document key should not be encrypted as the key should not leave the server. If the system is using server-side securing, the encrypted ticket data from the Cryptography module can be embedded within the document, and the document key can be used to encrypt the document (1840). If the system is using client-side securing, this is not needed.

[0153] The system can audit that a document was secured (1845). If the system is using server-side securing, the encrypted file can be returned to the client (1850). Otherwise the encrypted ticket data and the document key can be returned to the client (1850). If the system is using client-side securing, the document securer on the client can embed the encrypted ticket data and encrypt the document using the document key on the client (1855).

[0154] FIG. 19 is a block diagram illustrating server-side ACL evaluation workflow in the document control server of FIG. 9. When the server performs an operation that involves permissions, the server can first determine the authenticated user identity (1900). The encrypted server control information within the document can be decrypted (1910). The ticket in the encrypted control information can be used to retrieve the most recent document ACL from the storage service provider (1920). The Access Control Manager can evaluate the ACL, determining which permissions are relevant to the authenticated user (1930). The ACL may reference groups, and so the storage provider can be queried to determine which groups the authenticated user belongs to (1940).

[0155] FIG. 20 is a block diagram illustrating online document viewing workflow in the document control server of FIG. 9. Viewing a document while online can involve two major phases. The first phase involves determining which permissions the authenticated user has, and the second phase involves returning the document key to decrypt the document on the client. When a document is to be viewed online, a viewing application can open a secured document and recognize that the document is associated with the control server (e.g., the document can involve a security handler in the viewing client). Using the server RPC, the viewing application can transmit to the server the encrypted control information within the encrypt dictionary in the document (2000). The server can evaluate the ACL as an operation that involves permissions (2010), as described above in connection with FIG. 19. Then, the storage provider can be queried to ensure this document has not been revoked (2020). The document key can be extracted from the control information (2030). The server can audit the online viewing of this document (2040). The most recent ACL, the rules for viewing this document, as well as the document key can then be returned to the viewing client (2050). The viewing application can then enforce the permissions (e.g., the security handler can

inform the viewing application what permissions to enforce, and provide the decryption key such that the document can be viewed).

[0156] FIG. 21 is a block diagram illustrating revocation workflow in the document control server of FIG. 9. The client can send the encrypted control information to the server (2100). The server can determine whether the authenticated user has permission to revoke the document (2110), as described above in connection with FIG. 19. The server can then revoke the document (2120). The client can receive an acknowledgement (2130).

[0157] FIG. 22 is a block diagram illustrating audit events retrieval workflow in the document control server of FIG. 9. The client can send the encrypted control information to the server (2200). The server can determine whether the authenticated user has permission to get the audit history for this document (2210), as described above in connection with FIG. 19. The storage provider can be queried to determine what events are relevant to this document (2220). The client can then receive and display the audit information to the user (2230).

[0158] FIG. 23 is a block diagram illustrating a document control system with multiple document control servers 2360. The system can use a three tier architecture to provide reliability and scalability. Clients 2310, 2320, 2330 in an application tier 2300 communicate with the document control servers 2360 in a business logic tier 2350, which communicate with enterprise systems (e.g., a database management system (DBMS) 2380) in a storage tier 2370. All server state that is not specific to that particular instance of the server can be stored in the third tier 2370 so that multiple server instances can share such state.

[0159] When multiple document control server instances 2360 are used, requests can be routed to other servers if one goes down. A load balancer 2340 can handle routing of requests to the server instances 2360. Within a server itself, high reliability can be achieved by writing the server in a language using managed code, such as Java or a .NET language. In order to manage many canonical and non-canonical principals, two levels of cache can be provided for principal information. A server 2360 can have an in-memory cache of canonical mapping and group membership for recently queried canonical users. Many document control servers can share the secondary cache within the storage provider.

[0160] Should the desired information not exist within either of these caches, the servers can directly access the direct principal providers within the Access Control service provider and then cache the information both locally and within the storage provider. Group membership information should be batch processed such that it can be retrieved as needed in a reasonable amount of time. One of the document control servers, as a secondary service, can be designated a master and have the

responsibility of performing the batch processing tasks. In many cases, the actual securing can be done on the client to remove the overhead of transferring the document to and from the server and to reduce the load on the server. Likewise, with client-side securing, the client can also perform the document encryption, further decreasing server load.

[0161] The three-tier architecture allows server replicas to be added to scale to large enterprises. Documents can be tethered to a cluster of servers instead of to a specific hostname, as described above. DNS (Domain Naming System) round-robin can be added to the system to allow for additional hardware to act as document control servers. The servers can contain no state, so the hardware scalability concern can be reduced to the standard "one database" problem. Algorithms regarding principal management can be designed to be $O(1)$ for individual operations and $O(n)$ for aggregate operations (batch processing, etc.).

[0162] The invention and all of the functional operations described in this specification can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Apparatus of the invention can be implemented in a software product (e.g., a computer program product) tangibly embodied in a machine-readable storage device for execution by a programmable processor; and processing operations of the invention can be performed by a programmable processor executing a program of instructions to perform functions of the invention by operating on input data and generating output. The invention can be implemented advantageously in one or more software programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. Each software program can be implemented in a high-level procedural or object-oriented programming language, or in assembly or machine language if desired; and in any case, the language can be a compiled or interpreted language. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory, a random access memory and/or a machine-readable signal (e.g., a digital signal received through a network connection). Generally, a computer will include one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks, magneto-optical disks, and optical disks. Storage devices suitable for tangibly embodying software program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM (electrically programmable read-only memory), EEPROM (electrically erasable programmable read-only memory), and flash memory devices; magnetic disks such as

internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

[0163] The invention has been described in terms of particular embodiments. Other embodiments are within the scope of the following claims. For example, the operations of the invention can be performed in a different order and still achieve desirable results. The operations can be provided as a hosted service, using a subscription business model, and integrations can be performed with generally available system infrastructure available over the Internet. The document version control techniques can be implemented using peer-to-peer systems and techniques. Moreover, the sets of permissions for documents can be extended to cover various actions with respect to document content given different workflows (e.g., permissions that allow only certain people to sign a document, or portions of a document, and/or permissions that control who may fill out and/or view different sections of an electronic form).

[0164] Additionally, an alternative to always synchronizing policy updates but not necessarily other ACLs, can involve providing information regarding which ACLs in the system have changed. Synchronization operations can then be divided into high and low priority operations. High priority synchronizations can occur in the background more frequently, and provide indications of when information has changed. For example, an indication of which access control lists and policies have changed since the client's last synchronization. Low priority synchronization operations can entail how information has changed. For example, this can include the offline access information for every document in the system that has changed. Synchronizing how access control information has changed should be generally more resource intensive than a summary of what has changed. If access control for a document has been modified and the client is aware of a modification but has not performed a low priority synchronization, the system can be conservative and an implementation can prevent access to that document until the low priority synchronization has taken place.

Claims

1. A method comprising:

receiving a request from a client; and
pre-authorizing the client, in response to the request, to allow actions by a user as a member of a group of users by sending to the client of-line access information comprising a first key associated with the group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document.

2. The method of claim 1, wherein pre-authorizing the client comprises comparing a time of last recorded client-synchronization with a time of last change in user-group information for the user.
3. The method of one of the preceding claims, wherein pre-authorizing the client comprises comparing current user-group information for the user with received user-group information for the user from the client.
4. The method of one of the preceding claims, wherein the client allows actions with respect to the electronic document based on document-permissions information residing in the electronic document.
5. The method of one of the preceding claims, wherein the offline access information further comprises document-permissions information associated with multiple documents, including the electronic document, and the client allows actions with respect to the electronic document based on the document-permissions information.
6. The method of one of the preceding claims, wherein receiving a request comprises receiving a request from the client to take an action with respect to a second document.
7. The method of claim 6, further comprising verifying the user at the client as an authenticated user.
8. The method of claim 6 or 7, wherein the offline access information further comprises:
 - at least one user-specific key;
 - at least one group-specific key, including the first key; and
 - at least one set of document-permissions information associated with multiple documents.
9. The method of claim 8, further comprising receiving an offline audit log from the client.
10. The method of claim 8 or 9, wherein the at least one set of document-permissions information comprises one or more policies associated with the second document, and the offline access information further comprises a document revocation list.
11. The method of claim 8, 9 or 10, wherein the offline access information further comprises at least one set of document-permissions information associated with a specific document selected based on synchronization prioritization information.
12. A method comprising:

- receiving from a document control server, when online, offline access information comprising a first key associated with a group of users of the document control server; and
allowing access to an electronic document, when offline, by performing operations comprising using the first key to decrypt a second key in the electronic document and governing actions with respect to the electronic document based on document-permissions information associated with the electronic document. 5 10
13. The method of claim 12, wherein governing actions with respect to the electronic document comprises obtaining the document-permissions information from the electronic document. 15
14. The method of claim 12 or 13, wherein governing actions with respect to the electronic document comprises: 20
identifying a document policy reference in the electronic document; and
obtaining locally retained document-permissions information based on the document policy reference. 25
15. The method of one of the preceding claims 12 to 14, wherein the offline access information comprises at least one user-specific key, at least one group-specific key, including the first key, at least one set of document-permissions information associated with multiple documents, and a document revocation list. 30 35
16. The method of one of the preceding claims 12 to 15, further comprising preventing access to the document, when offline, if a difference between a current time and a receipt time of the offline access information exceeds a server-synchronization-frequency parameter. 40
17. The method of claim 16, wherein the server-synchronization-frequency parameter is specific to the document. 45
18. The method of one of the preceding claims 12 to 17, further comprising: 50
maintaining an offline audit log; and
uploading the offline audit log when online.
19. A method comprising: 55
encrypting an electronic document; and
incorporating into the encrypted electronic document an address of a document control server, document-permissions information, and an encryption key useable in decrypting the encrypted electronic document, the encryption key being encrypted with a key generated by, and associated with a group of users of, the document control server.
20. The method of claim 19, wherein the encryption key comprises a session key generated by the document control server, encrypting the electronic document comprises encrypting the electronic document using a document key, and incorporating comprises incorporating into the encrypted electronic document a document security payload comprising the document key and the document-permissions information, the document security payload being encrypted using the session key.
21. The method of claim 20, wherein the document security payload further comprises a document identifier assigned by the document control server, and incorporating further comprises incorporating into the encrypted electronic document a copy of the session key encrypted using a public key associated with the document control server.
22. The method of one of the preceding claims 19 to 21, wherein the document-permissions information specifies access permissions at a level of granularity smaller than the electronic document.
23. A software product tangibly embodied in a machine-readable medium, the software product comprising instructions operable to cause one or more data processing apparatus to perform operations comprising: 60
receiving a request from a client; and
pre-authorizing the client, in response to the request, to allow actions by a user as a member of a group of users by sending to the client offline access information comprising a first key associated with the group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document.
24. The software product of claim 23, wherein pre-authorizing the client comprises comparing a time of last recorded client-synchronization with a time of last change in user-group information for the user.
25. The software product of claim 23 or 24, wherein pre-authorizing the client comprises comparing current user-group information for the user with received user-group information for the user from the client.
26. The software product of one of the preceding claim 23 to 25, wherein the client allows actions with re-

spect to the electronic document based on document-permissions information residing in the electronic document.

27. The software product of one of the preceding claims 23 to 26, wherein the offline access information further comprises document-permissions information associated with multiple documents, including the electronic document, and the client allows actions with respect to the electronic document based on the document-permissions information. 5
28. The software product of one of the preceding claims 23 to 27, wherein receiving a request comprises receiving a request from the client to take an action with respect to a second document. 10
29. The software product of claim 28, wherein the operations further comprise verifying the user at the client as an authenticated user. 20
30. The software product of claim 28 or 29, wherein the offline access information further comprises:
 - at least one user-specific key; 25
 - at least one group-specific key, including the first key; and
 - at least one set of document-permissions information associated with multiple documents. 30
31. The software product of claim 30, wherein the operations further comprise receiving an offline audit log from the client.
32. The software product of claim 30 or 31, wherein the at least one set of document-permissions information comprises one or more policies associated with the second document, and the offline access information further comprises a document revocation list. 35
33. The software product of claim 30, 31 or 32, wherein the offline access information further comprises at least one set of document-permissions information associated with a specific document, and the operations further comprise selecting the document-specific document-permissions information based on synchronization prioritization information. 40
34. A software product tangibly embodied in a machine-readable medium, the software product comprising instructions operable to cause one or more data processing apparatus to perform operations comprising: 45
 - receiving from a document control server, when online, offline access information comprising a first key associated with a group of users of the 50

document control server; and
allowing access to an electronic document, when offline, by performing operations comprising using the first key to decrypt a second key in the electronic document and governing actions with respect to the electronic document based on document-permissions information associated with the electronic document.

35. The software product of claim 34, wherein governing actions with respect to the electronic document comprises obtaining the document-permissions information from the electronic document. 10
36. The software product of claim 34 or 35, wherein governing actions with respect to the electronic document comprises:
 - identifying a document policy reference in the electronic document; and
 - obtaining locally retained document-permissions information based on the document policy reference. 15
37. The software product of claim 34, 35 or 36, wherein the offline access information comprises at least one user-specific key, at least one group-specific key, including the first key, at least one set of document-permissions information associated with multiple documents, and a document revocation list. 20
38. The software product of one of the preceding claims 34 to 37, wherein the operations further comprise preventing access to the document, when offline, if a difference between a current time and a receipt time of the offline access information exceeds a server-synchronization-frequency parameter. 25
39. The software product of claim 38, wherein the server-synchronization-frequency parameter is specific to the document. 30
40. The software product of one of the preceding claims 34 to 39, wherein the operations further comprise:
 - maintaining an offline audit log; and
 - uploading the offline audit log when online. 35
41. A software product tangibly embodied in a machine-readable medium, the software product comprising instructions operable to cause one or more data processing apparatus to perform operations comprising: 40
 - encrypting an electronic document; and
 - incorporating into the encrypted electronic document an address of a document control server, 45

document-permissions information, and an encryption key useable in decrypting the encrypted electronic document, the encryption key being encrypted with a key generated by, and associated with a group of users of, the document control server.

42. The software product of claim 41, wherein the encryption key comprises a session key generated by the document control server, encrypting the electronic document comprises encrypting the electronic document using a document key, and incorporating comprises incorporating into the encrypted electronic document a document security payload comprising the document key and the document-permissions information, the document security payload being encrypted using the session key.

43. The software product of claim 42, wherein the document security payload further comprises a document identifier assigned by the document control server, and incorporating further comprises incorporating into the encrypted electronic document a copy of the session key encrypted using a public key associated with the document control server.

44. The software product of claim 41, 42 or 43, wherein the document-permissions information specifies access permissions at a level of granularity smaller than the electronic document.

45. A system comprising:

a document control server that synchronizes offline access information with a client in response to a client request, the offline access information comprising a first key associated with a group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document; and
the client that allows access to the electronic document, when offline, by a user as a member of the group, using the first key to decrypt the second key in the electronic document and governing actions with respect to the electronic document based on document-permissions information associated with the electronic document.

46. The system of claim 45, wherein the electronic document comprises the document-permissions information.

47. The system of claim 46, wherein the second key comprises a session key generated by the document control server, and the electronic document further comprises a document security payload

comprising a document key and the document-permissions information, the document security payload being encrypted using the session key.

48. The system of one of the preceding claims 45 to 47, wherein the offline access information further comprises:

at least one user-specific key;
at least one group-specific key, including the first key; and
at least one set of document-permissions information associated with multiple documents.

49. The system of one of the preceding claims 45 to 48, wherein the client comprises an agent that periodically contacts the document control server to synchronize the offline access information.

50. The system of one of the preceding claims 45 to 49, wherein the document control server comprises:

a server core with configuration and logging components;
an internal services component that provides functionality across dynamically loaded methods; and
dynamically loaded external service providers, including one or more access control service providers.

51. The system of one of the preceding claims 45 to 50, further comprising:

a business logic tier comprising a cluster of document control servers, including the document control server;
an application tier including the client comprising a viewer client, a securing client, and an administration client; and
a load balancer that routes client requests to the document control servers.

52. The system of one of the preceding claims 45 to 51, wherein the client request comprises a request from the client to take an action with respect to a second document.

53. The system of claim 52, wherein the document control server comprises a permissions-broker server including a translation component, the second document comprises a document secured previously by the permissions-broker server, and the translation component being operable to translate first document-permissions information in a first permissions-definition format into second document-permissions information in a second permissions-definition format in response to the request being re-

ceived from the client.

54. The system of claim 52 or 53, wherein the server comprises a permissions-broker server operable to identify information associated with the second document in response to the request, the associated information being retained at the server and indicating a third electronic document different from and associated with the second document, the server being operable to relate information concerning the third electronic document to the client to facilitate the action to be taken.

55. The system of one of the preceding claims 52 to 54, wherein the server comprises a permissions-broker server operable to obtain and send, in response to the request, a software program comprising instructions operable to cause one or more data processing apparatus to perform operations effecting an authentication procedure, and the client uses the authentication program to identify a current user and control the action with respect to the second document based on the current user and document-permissions information associated with the second document.

56. A system comprising:

server means for transparently providing offline access information for controlled documents to pre-authorize a client to allow actions by a user as a member of a group of users, the offline access information comprising a first key associated with the group, the first key being useable at the client to access an electronic document by decrypting a second key in the electronic document; and
client means for accessing the electronic document using the offline access information.

57. The system of claim 56, further comprising:

server means for dynamically obtaining and sending authentication processes in response to client requests to take actions with respect to electronic documents; and
client means for interfacing with a received authentication process to identify a current user and for controlling actions with respect to electronic documents based on the current user and document-permissions information.




FIG. 1




FIG. 2




FIG. 3

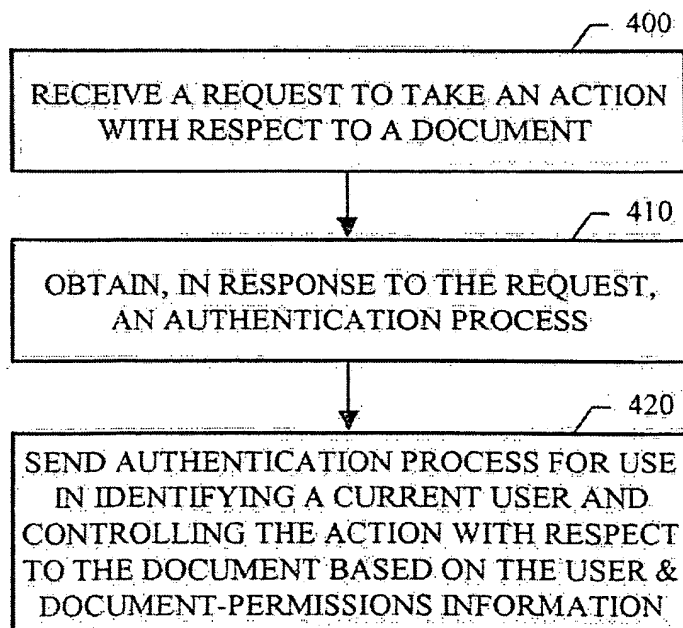


FIG. 4




FIG. 5

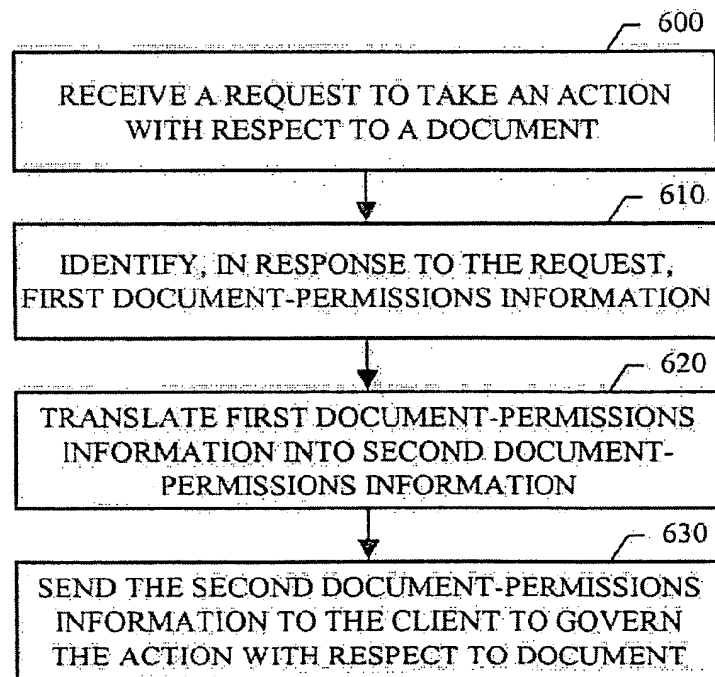


FIG. 6




FIG. 7




FIG. 8




FIG. 9

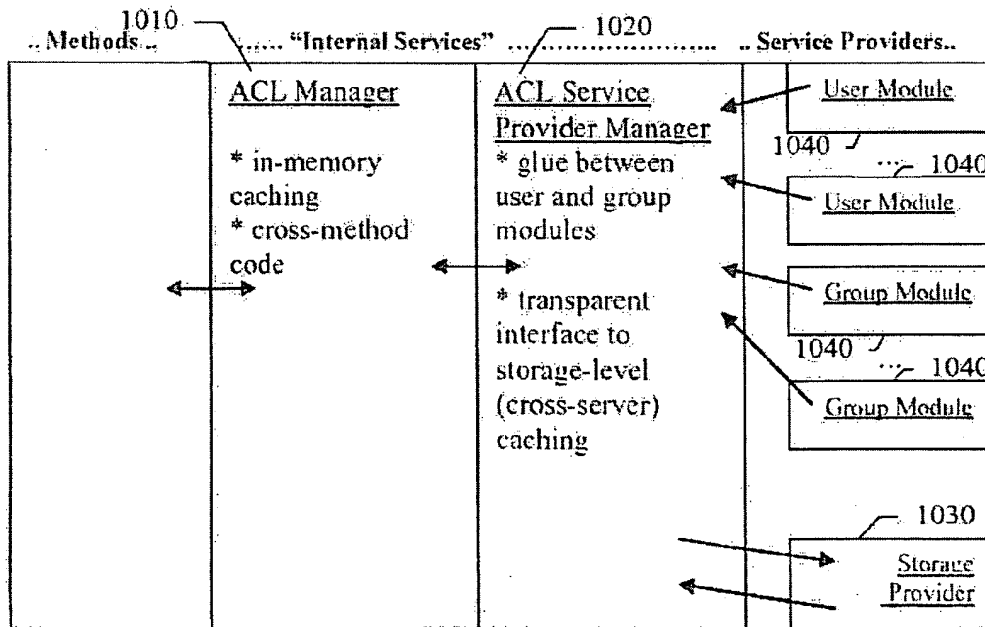
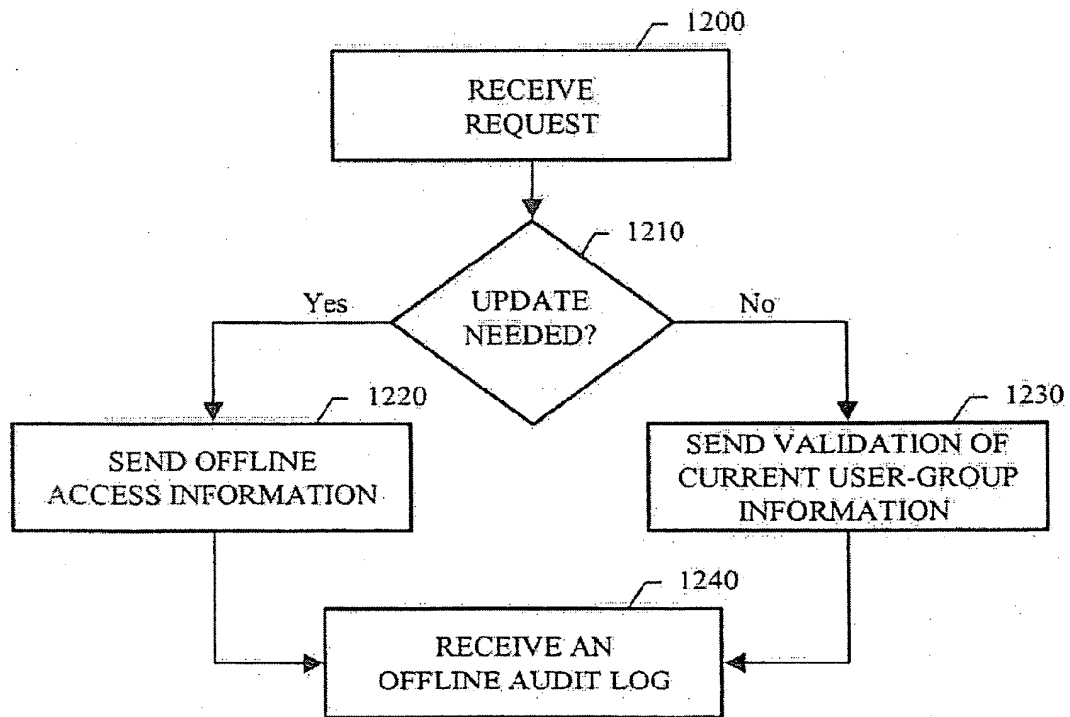



FIG. 10



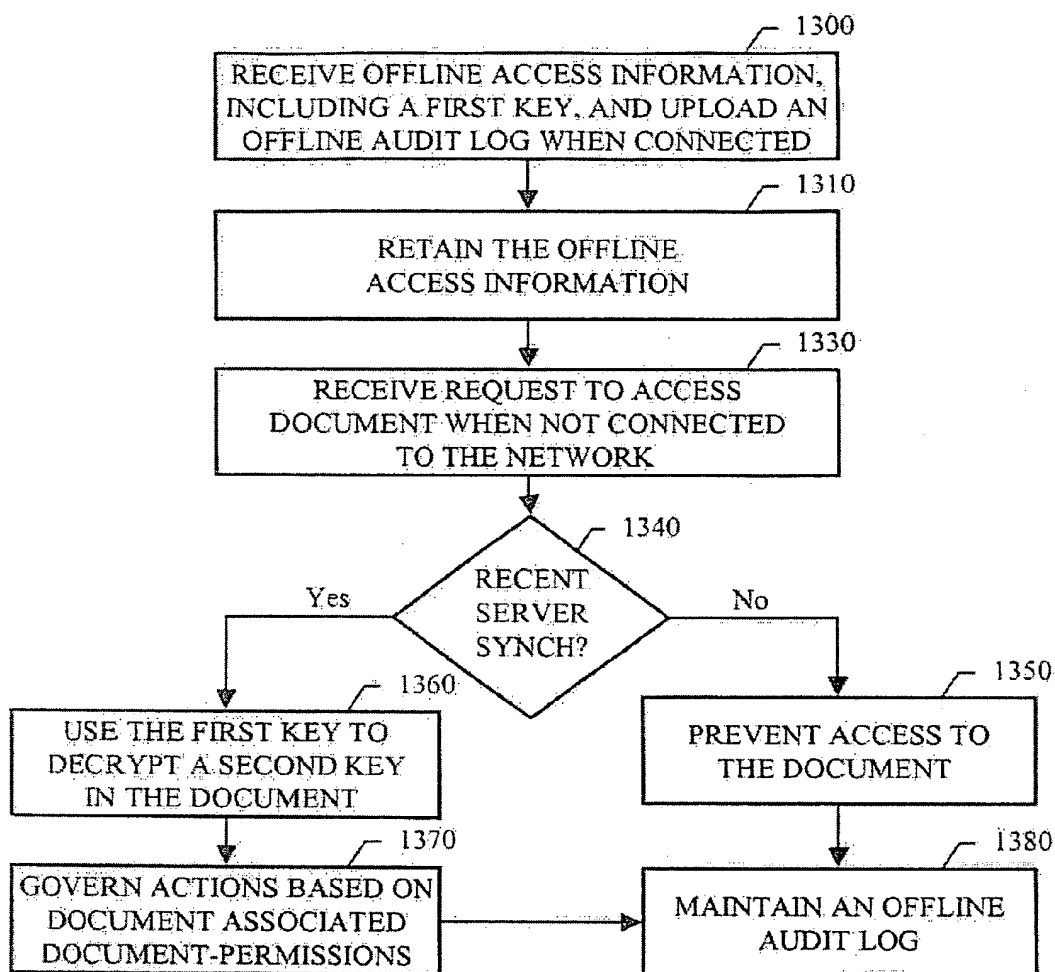


FIG. 13




FIG. 14

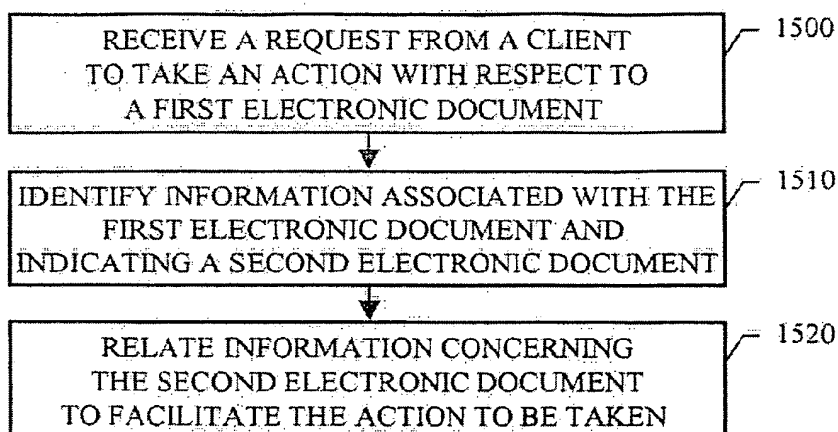


FIG. 15




FIG. 16

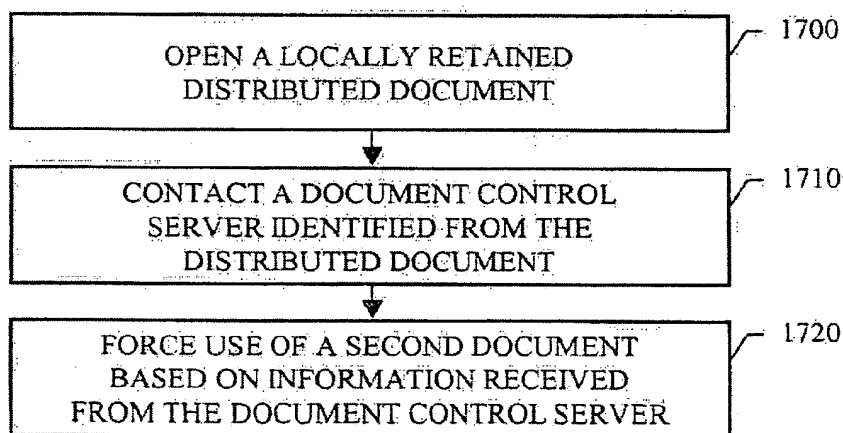
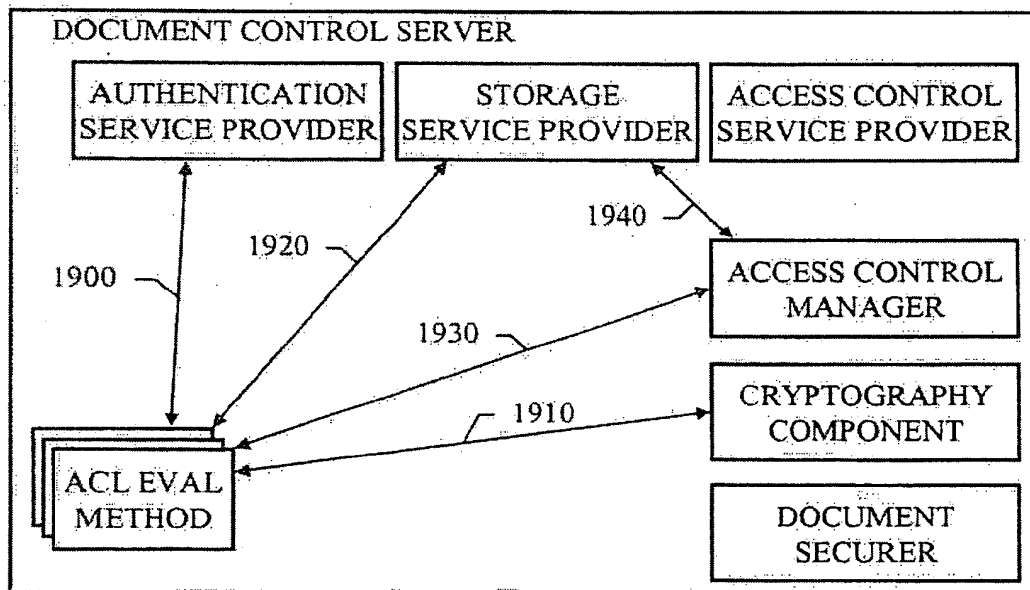
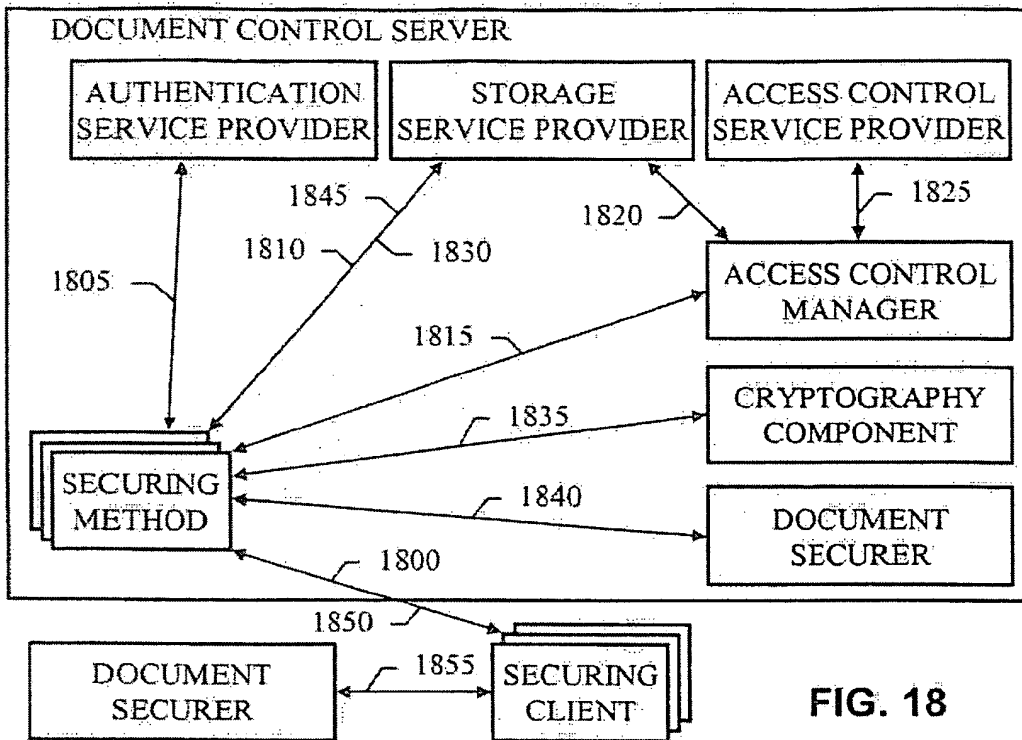
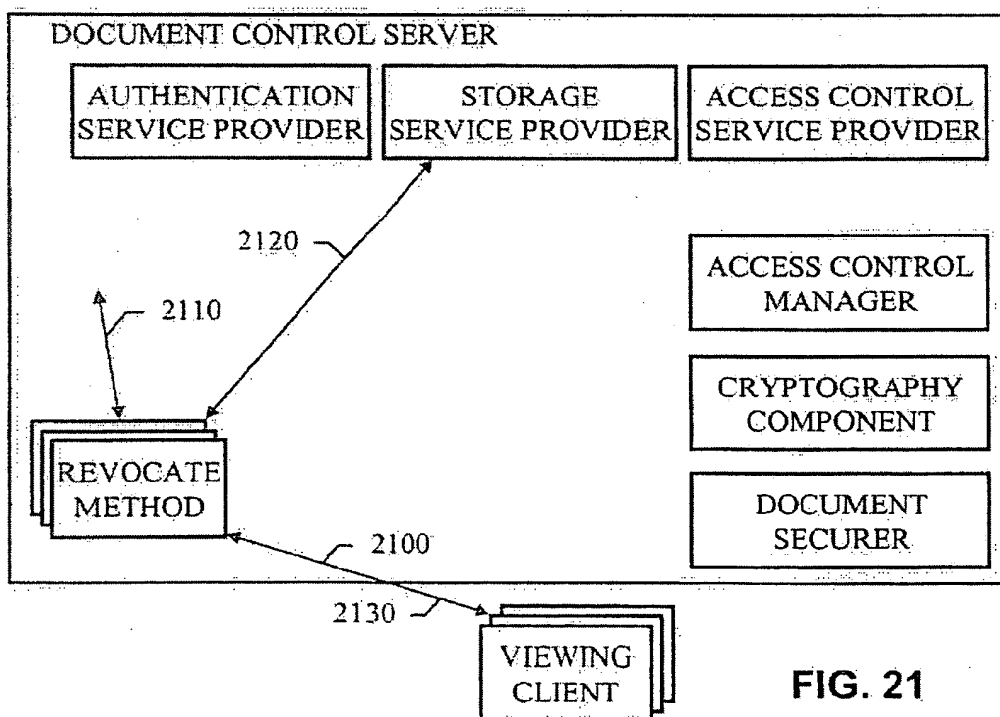
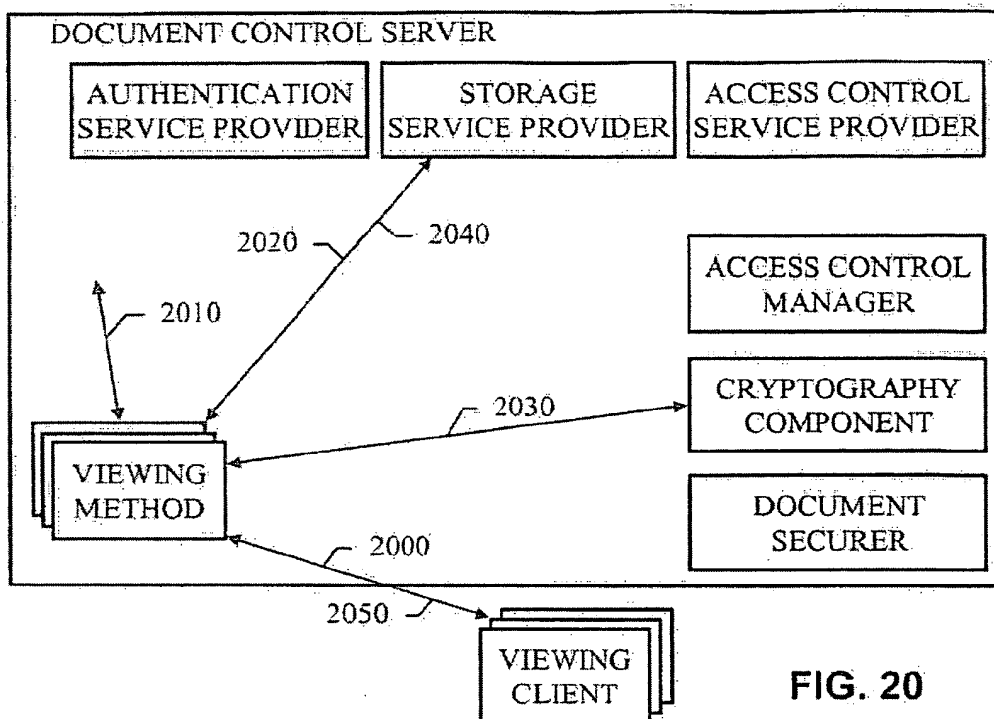


FIG. 17





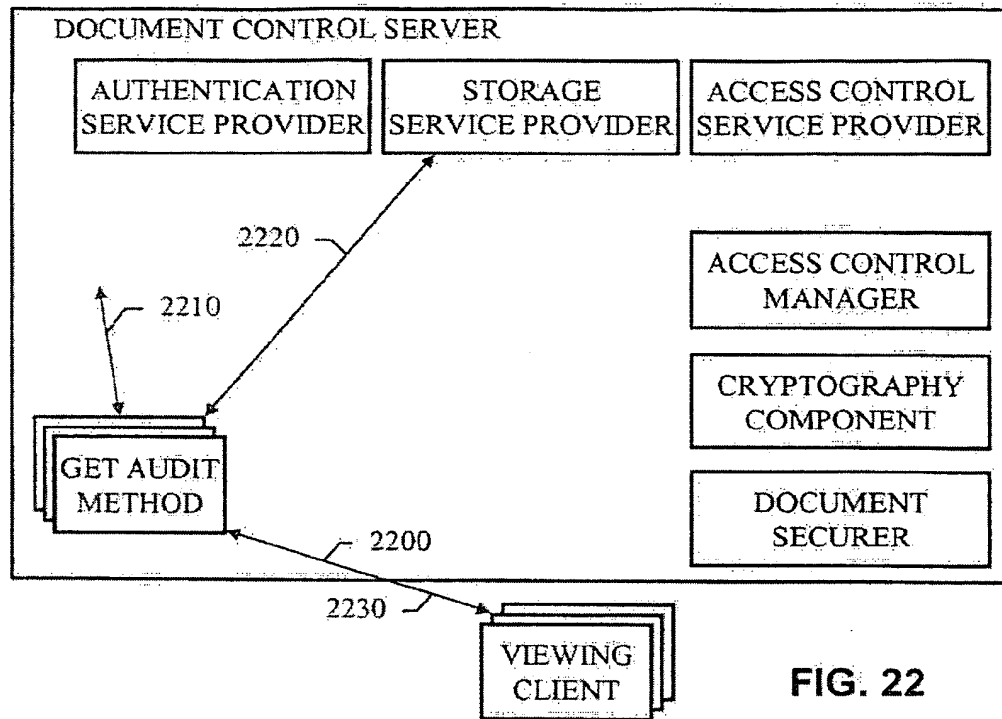


FIG. 22




FIG. 23



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 04 02 5826

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	US 5 757 907 A (COOPER ET AL) 26 May 1998 (1998-05-26)	1-7, 10-14, 19-24, 26-29, 31-36, 38,39, 41-47, 49-57	G06F1/00
Y	* figure 5 *	8,9, 15-18, 25,30, 37,40,48	
	* column 2, line 25 - column 5, line 32 * * column 9, line 1 - column 14, line 20 * -----		
X	EP 1 307 000 A (SONY CORPORATION) 2 May 2003 (2003-05-02)	1-7, 10-14, 19-24, 26-29, 32-36, 38,39, 41-47, 49-57	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
Y	* abstract *	8,9, 15-18, 25,30, 31,37, 40,48	G06F H04L
	* paragraph [0004] - paragraph [0011] * * paragraph [0033] - paragraph [0064] * ----- -/--		
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 22 February 2005	Examiner Anticoli, C
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

3
EPO FORM 1503 03.82 (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 04 02 5826

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
Y A	US 5 613 089 A (HORNBUCKLE ET AL) 18 March 1997 (1997-03-18) * column 2, line 50 - column 4, line 56 *	9,16-18, 25,31,40 1-8, 10-15, 19-24, 26-30, 32-39, 41-57	
Y A	----- EP 1 326 157 A (PERVASIVE SECURITY SYSTEMS INC) 9 July 2003 (2003-07-09) * paragraph [0009] - paragraph [0014] * * paragraph [0025] - paragraph [0038] * -----	8,15,30, 37,48 1-7, 9-14, 16-29, 31-36, 38-47, 49-57	
The present search report has been drawn up for all claims			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
Place of search Munich		Date of completion of the search 22 February 2005	Examiner Anticoli, C
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

3

EPC FORM 1503 03.82 (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 04 02 5826

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

22-02-2005

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5757907	A	26-05-1998	BR 9501522 A	21-11-1995
			CA 2145926 A1	26-10-1995
			DE 69531079 D1	24-07-2003
			DE 69531079 T2	22-04-2004
			EP 0679980 A1	02-11-1995
			JP 7295801 A	10-11-1995
			KR 200443 B1	15-06-1999

EP 1307000	A	02-05-2003	EP 1307000 A1	02-05-2003
			CN 1463514 T	24-12-2003
			WO 02080447 A1	10-10-2002
			TW 561753 B	11-11-2003
			US 2003177393 A1	18-09-2003

US 5613089	A	18-03-1997	US 5497479 A	05-03-1996
			AT 143511 T	15-10-1996
			AU 641397 B2	23-09-1993
			AU 5646490 A	29-11-1990
			CN 1048271 A ,C	02-01-1991
			CN 1208198 A	17-02-1999
			DE 69028705 D1	31-10-1996
			DE 69028705 T2	06-03-1997
			EP 0478571 A1	08-04-1992
			JP 4504794 T	20-08-1992
			WO 9013865 A1	15-11-1990
			CA 2053261 A1	29-10-1990
			US 5388211 A	07-02-1995
			US 5649187 A	15-07-1997

EP 1326157	A	09-07-2003	EP 1326156 A2	09-07-2003
			EP 1320010 A2	18-06-2003
			EP 1320011 A2	18-06-2003
			EP 1320012 A2	18-06-2003
			EP 1320013 A2	18-06-2003
			EP 1320014 A2	18-06-2003
			EP 1320015 A2	18-06-2003
			EP 1326157 A2	09-07-2003
			EP 1320016 A2	18-06-2003
			EP 1320017 A2	18-06-2003
			EP 1320018 A2	18-06-2003
			EP 1324565 A1	02-07-2003
			JP 2003228519 A	15-08-2003
			JP 2003228520 A	15-08-2003
			JP 2003223353 A	08-08-2003
			JP 2003218851 A	31-07-2003
			JP 2003248658 A	05-09-2003

EPO FORM P0455

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 04 02 5826

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

22-02-2005

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1326157	A	JP 2003242015 A	29-08-2003
		US 2004103202 A1	27-05-2004
		US 2003110131 A1	12-06-2003
		US 2003120601 A1	26-06-2003
		US 2003110169 A1	12-06-2003
		US 2003110397 A1	12-06-2003
		US 2003120684 A1	26-06-2003
		US 2003154381 A1	14-08-2003

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82